Transversal structures on triangulations: a combinatorial study and straight-line drawings

Éric Fusy

Algorithms Project, INRIA Rocquencourt and LIX, École Polytechnique

Abstract

This article focuses on a combinatorial structure specific to triangulated plane graphs with quadrangular outer face and no separating triangle, which are called irreducible triangulations. The structure has been introduced by Xin He under the name of regular edge-labelling and consists of two bipolar orientations that are transversal. For this reason, the terminology used here is that of transversal structures. The main results obtained in the article are a bijection between irreducible triangulations and ternary trees, and a straight-line drawing algorithm for irreducible triangulations. For a random irreducible triangulation with n vertices, the grid size of the drawing is asymptotically with high probability $11n/27 \times 11n/27$ up to an additive error of $\mathcal{O}(\sqrt{n})$. In contrast, the best previously known algorithm for these triangulations only guarantees a grid size $(\lceil n/2 \rceil - 1) \times \lfloor n/2 \rfloor$.

Key words: triangulations, bipolar orientations, bijection, straight-line drawing
PACS: 02.10.E

1 Introduction

A *plane graph*, or planar map, is a connected graph embedded in the plane without edge-crossings, and considered up to orientation-preserving homeomorphism. Many drawing algorithms for plane graphs [1,4,9,12,18,20,27] endow the graph with a particular structure, from which coordinates are assigned to vertices in a natural way. For example, triangulations, i.e., plane graphs with only triangular faces, are characterized by the fact that their inner edges can be partitioned into three spanning trees with specific incidence relations, the

---

Email address: `eric.fusy@inria.fr` (Éric Fusy).

so-called Schnyder woods [27]. These spanning trees yield a natural method to assign coordinates to a vertex, by counting the number of faces in specific regions. Placing the vertices accordingly on an integer grid and linking adjacent vertices by segments yields a straight-line drawing algorithm, which can be refined to produce a drawing on an integer grid $(n-2) \times (n-2)$, with n the number of vertices, see [4,28].

In this article we focus on so-called *irreducible triangulations*, which are plane graphs with quadrangular outer face, triangular inner faces, and no separating triangle, i.e., each 3-cycle is the boundary of a face. Irreducible triangulations form an important class of triangulations, as they are closely related to 4-connected triangulations. In addition, as discussed in [3], many plane graphs, including bipartite plane graphs and 4-connected plane graphs with at least 4 outer vertices, can be triangulated into an irreducible triangulation. There exist more compact straight-line drawing algorithms for irreducible triangulations [19,22], the size of the grid being guaranteed to be $(\lceil n/2 \rceil - 1) \times \lfloor n/2 \rfloor$ in the worst case. By investigating a bijection with ternary trees, we have observed that each irreducible triangulation T can be endowed with a so-called *transversal structure*, which can be summarized as follows. Calling W, N, E, S (like West, North, East, South) the four outer vertices of T in clockwise order, the inner edges of T can be oriented and partitioned into two sets: red edges that "flow" from S to N, and blue edges that "flow" from W to E. As we learned after completing a first draft of this paper, X. He [18] has defined the same structure under the name of regular edge-labelling, and derived a nice algorithm of rectangular-dual drawing, which has been recently applied to the theory of cartograms [8,29]. We give two equivalent definitions of transversal structures in Section 2: one without orientations called a transversal edge-partition, and one with orientations called a transversal pair of bipolar orientations (which corresponds to the regular edge-labelling of X. He). Transversal structures characterize irreducible triangulations in the same way as Schnyder Woods characterize triangulations, and they share similar combinatorial properties. In particular, we show in Section 3 that the set of transversal structures of an irreducible triangulation is a distributive lattice, and that the "flip" operation has a simple geometric interpretation (Theorem 2).
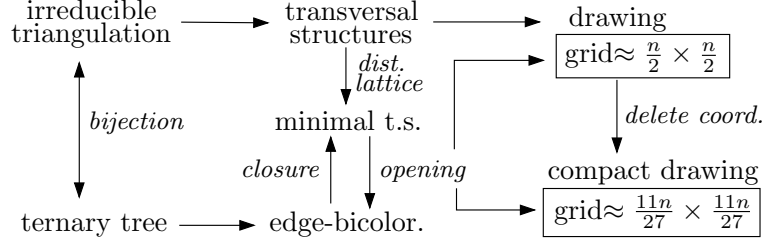
The transversal structure at the bottom of the lattice, called *minimal*, has a strong combinatorial role, as it allows us to establish a bijection between ternary trees and irreducible triangulations. The bijection, called *closure mapping*, is described in Section 4. The mapping from ternary trees to irreducible triangulations relies on "closure operations", as introduced by G. Schaeffer in his PhD [25], see also [17,24]. This bijection has in fact brought about our discovery of transversal edge-partitions, as a natural *edge-bicoloration* of a ternary tree is mapped to the minimal transversal edge-partition of the associated triangulation (similar, the bijection of [24] maps the structure

of Schnyder woods). Classical algorithmic applications of bijections between trees and plane graphs are random generation —as detailed in [25,26]— and encoding algorithms for plane graphs —as detailed in [17,24]— with application to mesh compression in computational geometry. The closure mapping presented in this article yields linear time procedures for random sampling (under a fixed-size uniform distribution) and optimal encoding (in the information theoretic sense) of *4-connected triangulations*. These algorithms are described in the thesis of the author [16]. The focus in this article (besides the graph drawing algorithm) is on the application to counting; the bijection yields a combinatorial way to enumerate rooted 4-connected triangulations, which were already counted by Tutte in [30] using algebraic methods.

In Section 5, we derive from transversal structures a straight-line drawing algorithm for irreducible triangulations. In a similar way as algorithms using Schnyder Woods [4,27], the drawing is obtained by using face-counting operations. Our algorithm outputs a straight line drawing on an integer grid of half-perimeter $n - 1$ if the triangulation has n vertices (Theorem 5). This is to be compared with previous algorithms for irreducible triangulations by He [19] and Miura *et al* [22]. The latter produces a grid $(\lceil n/2 \rceil - 1) \times \lfloor n/2 \rfloor$; the half-perimeter is also $n - 1$, but the aspect ratio of the outer face is better. However, the algorithms of [19] and [22] rely on a particular order to treat the vertices, called *canonical ordering*, and a step of coordinate-shifting makes them difficult to implement and to carry out by hand. In contrast, our algorithm can readily be performed on a piece of paper, because the coordinates of the vertices are computed independently with simple face-counting operations.

Furthermore, some coordinate-deletions can be performed on the drawing obtained using the face-counting algorithm, with the effect of reducing the size of the grid (Theorem 6). For an irreducible triangulation with n vertices taken uniformly at random and endowed with its minimal transversal structure (for the distributive lattice), we show in Section 6 (Theorem 7) that the size of the grid after coordinate-deletions is asymptotically with high probability $11n/27 \times 11n/27$ up to an additive error of order $\sqrt{n}$. Compared to [19] and [22], we do not improve on the size of the grid in the worst case, but we improve asymptotically with high probability by a reduction-factor $27/22$ on the width and height of the grid, see Figure 12 for an example. The proof of the grid size $11n/27 \times 11n/27$ makes use of several ingredients: a combinatorial interpretation of coordinate-deletions, the bijection with ternary trees, and modern tools of analytic combinatorics such as the quasi-power theorem [14].

The following diagram summarizes the connections between the different combinatorial structures and the role they play for the drawing algorithms.

3

$$\text{irreducible triangulation} \xrightarrow{\hspace{2cm}} \text{transversal structures} \xrightarrow{\hspace{2cm}} \text{drawing} \quad \boxed{\text{grid} \approx \tfrac{n}{2} \times \tfrac{n}{2}}$$

$$\Big\updownarrow \text{bijection} \qquad \Big\downarrow \substack{\text{dist.} \\ \text{lattice}} \qquad \text{minimal t.s.}$$

$$\text{ternary tree} \xrightarrow{\hspace{2cm}} \text{edge-bicolor.} \qquad \text{compact drawing} \quad \boxed{\text{grid} \approx \tfrac{11n}{27} \times \tfrac{11n}{27}}$$

2   Transversal structures: definitions

In this section we give two definitions of transversal structures, one without orientations called *transversal edge-partition* and one with orientations called *transversal pair of bipolar orientations*. As we will prove in Proposition 2, the two definitions are in fact equivalent, i.e., the additional information given by the orientations is redundant. The definition without orientations is more convenient for the combinatorial study (lattice property and bijection with ternary trees), while the definition with orientations fits better to describe the straight-line drawing algorithm in Section 5.
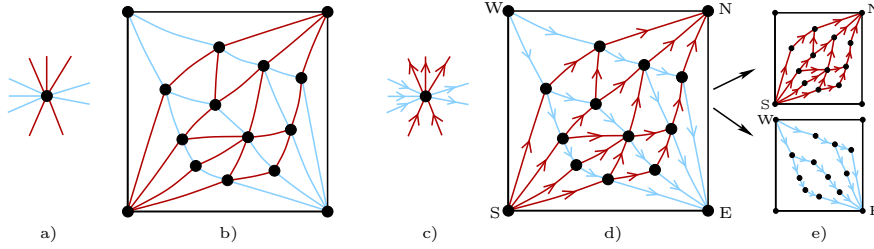

Fig. 1. Transversal edge-partition: local condition (a) and a complete example (b). In parallel, transversal pair of bipolar orientations: local condition (c) and a complete example (d) with the two induced bipolar orientations (e).

2.1   Transversal edge-partition

Let T be an irreducible triangulation. Edges and vertices of T are said to be *inner* or *outer* whether they are incident to the outer face or not. A *transversal edge-partition* of T is a combinatorial structure where each inner edge of T is given a color —red or blue— such that the following conditions are satisfied.

C1 (Inner vertices): In clockwise order around each inner vertex v, the edges incident to v form: a non empty interval of red edges, a non empty interval of blue edges, a non empty interval of red edges, and a non empty interval of blue edges, see Figure 1(a).

4

C2 (Outer vertices): Writing a_1, a_2, a_3, a_4 for the outer vertices of T in clockwise order, all inner edges incident to a_1 and to a_3 are of one color and all inner edges incident to a_2 and to a_4 are of the other color.

An example of transversal edge-partition is illustrated in Figure 1(b), where red edges are DARK GREY and blue edges are LIGHT GREY (the same convention will be used for all figures).

2.2 Transversal pair of bipolar orientations

An orientation of a graph G is said to be *acyclic* if it has no oriented circuit (a circuit is an oriented simple cycle of edges). Given an acyclic orientation of G, a vertex having no ingoing edge is called a *source*, and a vertex having no outgoing edge is called a *sink*. A *bipolar orientation* is an acyclic orientation with a unique source, denoted s, and a unique sink, denoted t. Such an orientation is also characterized by the fact that, for each vertex $v \neq \{s, t\}$, there exists an oriented path from s to t passing by v, see [10] for a detailed discussion. An important property of a bipolar orientation on a *plane graph* is that the edges incident to a vertex $v \neq \{s, t\}$ are partitioned into a non-empty interval of ingoing edges and a non-empty interval of outgoing edges; and, dually, each face f of M has two particular vertices s_f and t_f such that the boundary of f consists of two non-empty oriented paths both going from s_f to t_f, the path with f on its right (left) is called the *left lateral path* (*right lateral path*, resp.) of f.

Let T be an irreducible triangulation. Call N, E, S and W the four outer vertices of T in clockwise order around the outer face. A *transversal pair of bipolar orientations* of T is a combinatorial structure where each inner edge of T is given a direction and a color —red or blue— such that the following conditions are satisfied (see Figure 1(d) for an example):

C1' (Inner vertices): In clockwise order around each inner vertex v of T, the edges incident to v form: a non empty interval of outgoing red edges, a non empty interval of outgoing blue edges, a non empty interval of ingoing red edges, and a non empty interval of ingoing blue edges, see Figure 1(c).
C2' (Outer vertices): All inner edges incident to N, E, S and W are respectively ingoing red, ingoing blue, outgoing red, and outgoing blue.

This structure is also considered in [18,20] under the name of *regular edge labelling*.

Proposition 1 *The orientation of the edges given by a transversal pair of bipolar orientation is acyclic. The sources are W and S, and the sinks are E and N.*

PROOF. Let T be an irreducible triangulation endowed with a transversal pair of bipolar orientations. Assume there exists a circuit, and consider a minimal one $\mathcal{C}$, i.e., the interior of $\mathcal{C}$ is not included in the interior of any other circuit. It is easy to check from Condition C2' that $\mathcal{C}$ is not the boundary of a face of T. Thus the interior of $\mathcal{C}$ contains at least one edge e incident to a vertex v of $\mathcal{C}$. Assume that e is going out of v. Starting from e, it is always possible, when reaching a vertex inside $\mathcal{C}$, to go out of that vertex toward one of its neighbours. Indeed, a vertex inside $\mathcal{C}$ is an inner vertex of T, hence has positive outdegree according to Condition C1'. Thus, there is an oriented path starting from e, that either loops $\mathcal{C}$ into a circuit in the interior of $\mathcal{C}$ — impossible by minimality of $\mathcal{C}$— or reaches $\mathcal{C}$ again —impossible as a chordal path for $\mathcal{C}$ would yield two smaller circuits. Thus, the orientation is acyclic. Finally Condition C1' ensures that no inner vertex of T can be a source or a sink, and Condition C2' ensures that W and S are sources, and E and N are sinks. $\qquad\square$

The following corollary of Proposition 1, also proved in [18], explains the name of transversal pair of bipolar orientations, see also Figure 1(e).

Corollary 1 ([18]) *Let T be an irreducible triangulation endowed with a transversal pair of bipolar orientation. Then the (oriented) red edges induce a bipolar orientation of the plane graph obtained from T by removing W, E, and all non red edges. Similarly, the blue edges induce a bipolar orientation of the plane graph obtained from T by deleting S, N, and all non blue edges.*

3 Lattice property of transversal edge-partitions

We investigate on the set $\mathcal{E}(T)$ of transversal edge-partitions of a fixed irreducible triangulation T. Kant and He [20] have shown that $\mathcal{E}(T)$ is not empty and that an element of $\mathcal{E}(T)$ can be computed in linear time. In this section, we prove that $\mathcal{E}(T)$ is a distributive lattice. This property is to be compared with the lattice property of other similar combinatorial structures on plane graphs, such as bipolar orientations [23] and Schnyder woods [6,13]. Our proof takes advantage of the property that the set of orientations of a plane graph with a prescribed outdegree for each vertex is a distributive lattice. To apply this result to transversal edge-partitions, we establish a bijection between $\mathcal{E}(T)$ and some orientations with prescribed vertex outdegrees on an associated plane graph, called the *angular graph* of T.

3.1 Lattice structure of α-orientations on a plane graph

Let us first recall the definition of a distributive lattice. A *lattice* is a partially ordered set $(E, \leq)$ such that, for each pair (x, y) of elements of E, there exists a unique element $x \wedge y$ and a unique element $x \vee y$ satisfying the conditions:

- $x \wedge y \leq x$, $x \wedge y \leq y$, and $\forall z \in E$, $z \leq x$ and $z \leq y$ implies $z \leq x \wedge y$,
- $x \vee y \geq x$, $x \vee y \geq y$, and $\forall z \in E$, $z \geq x$ and $z \geq y$ implies $z \geq x \vee y$.

In other words, each pair admits a unique common lower element dominating all other common lower elements, and the same holds with common upper elements. The lattice is said to be *distributive* if the operators $\wedge$ and $\vee$ are distributive with respect to each other, i.e., $\forall(x, y, z) \in E$, $x \wedge (y \vee z) = (x \wedge y) \vee (x \wedge z)$ and $x \vee (y \wedge z) = (x \vee y) \wedge (x \vee z)$. The nice feature of distributive lattices is that, in most cases, moving from an element of the lattice to a covering lower (upper) element has a simple geometric interpretation, which we informally call a flip (flop, respectively). As we recall next, in the case of orientations of a plane graph with prescribed vertex outdegrees, the flip (flop) operation consists in reversing a clockwise circuit (counter-clockwise circuit, respectively).

Given a plane graph $G = (V, E)$, and a function $\alpha : V \to \mathbf{N}$, an α-*orientation* of G is an orientation of G such that each vertex v of G has outdegree $\alpha(v)$. An oriented circuit $\mathcal{C}$ is called *essential* if it has no chordal path (a *chordal path* is an oriented path of edges in the interior of $\mathcal{C}$ with the two extremities on $\mathcal{C}$).

Theorem 1 (Ossona de Mendez [23], Felsner [13]) *Given $G = (V, E)$ a plane graph and $\alpha : V \to \mathbf{N}$ a function, the set of α-orientations of G is either empty or is a distributive lattice. The flip operation consists in reversing the orientation of an essential clockwise circuit.*

3.2 Bijection with orientations of the angular graph

Let T be an irreducible triangulation. The *angular graph* of T is the bipartite plane graph Q with vertex set V_Q consisting of vertices and inner faces of T, and edge set corresponding to the incidences between these vertices and faces, see Figure 2. In the sequel, vertices of Q corresponding to vertices (inner faces) of T are black (white, respectively), and the whole vertex set of Q is denoted V_Q. The terminology is due to the fact that each edge of Q corresponds to an angle (v, f) of T (the angles incident to the outer face of T are not considered). Notice also that each inner face of $Q(T)$ is quadrangular.
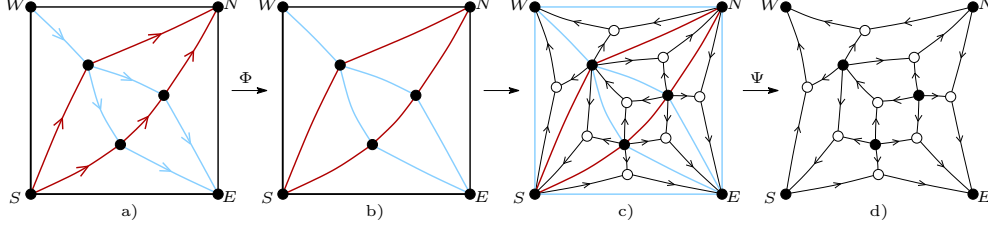
Fig. 2. Given an irreducible triangulation T endowed with a transversal pair of bipolar orientations Z (Fig.a) and the induced transversal edge-partition $E = \Phi(Z)$ (Fig.b), construction of the angular graph $Q(T)$ and of the α_0-orientation of $Q(T)$ image of E by the mapping Ψ (Fig.c and d).

We consider the function $\alpha_0 : V_Q \to \mathbf{N}$ specified as follows,

- each black vertex v of $Q(T)$ corresponding to an inner vertex of T satisfies $\alpha_0(v) = 4$,
- each white vertex v of $Q(T)$ satisfies $\alpha_0(v) = 1$,
- the outer black vertices satisfy $\alpha_0(N) = \alpha_0(S) = 2$, $\alpha_0(W) = \alpha_0(E) = 0$.

Proposition 2 *Given T an irreducible triangulation and $Q(T)$ the angular graph of T, the following sets are in bijection:*

- *transversal pairs of bipolar orientations of T,*
- *transversal edge-partitions of T,*
- *α_0-orientations of $Q(T)$.*

In particular, Proposition 2 ensures that the definitions of transversal edge-partition and of transversal pair of bipolar orientations are equivalent, i.e., the orientation of edges is a information once the colors are fixed.

Given T an irreducible triangulation and $Q(T)$ the angular graph of T, we denote by $\mathcal{B}$ the set of transversal pairs of bipolar orientations of T, $\mathcal{E}$ the set of transversal edge-partitions of T, and $\mathcal{O}$ the set of α_0-orientations of $Q(T)$. To prove Proposition 2, we introduce two mappings Φ and Ψ respectively from $\mathcal{B}$ to $\mathcal{E}$ and from $\mathcal{E}$ to $\mathcal{O}$.

Given $Z \in \mathcal{B}$, $\Phi(Z)$ is simply the edge-bicoloration induced by Z, as illustrated in Figure 2(a)-(b). It is straightforward that $\Phi(Z) \in \mathcal{E}$. In addition, Φ is clearly injective. Indeed, starting from a transversal edge-partition, the directions of edges of $Q(T)$ are assigned greedily so as to satisfy the local rules of a transversal pair of bipolar orientations. The fact that the propagation of edge directions is done without conflict is not straightforward; it will follow from the surjectivity of the mapping Φ, to be proved later.

Given $X \in \mathcal{E}$, we define $\Psi(X)$ as the following orientation of $Q(T)$. First, color blue the four outer edges of T. Then, for each angle (v, f) of T, orient the corresponding edge of $Q(T)$ out of v if (v, f) is bicolored, and toward v if

(v, f) is unicolored. (An angle (v, f) of T is called *bicolored* if it is delimited by two edges of T having different colors, and is called unicolored otherwise.) Condition C1 implies that all inner black vertices of $Q(T)$ have outdegree 4. In addition, Condition C2 and the fact that the four outer edges of T have been colored blue imply that E and W have outdegree 0 and that N and S have outdegree 2.

The following lemma ensures that all white vertices have outdegree 1 in $\Psi(X)$, so that $\Psi(X)$ is an α_0-orientation.

Lemma 1 *Let T be a plane graph with quadrangular outer face, triangular inner faces, and endowed with a transversal edge-partition, the four outer edges being additionally colored blue. Then there is no mono-colored inner face, i.e., each inner face of T has two sides of one color and one side of the other color.*

PROOF. Let Λ be the number of bicolored angles of T and let n be the number of inner vertices of T. Condition C1 implies that there are $4n$ bicolored angle incident to an inner vertex of T. Condition C2 and the fact that all outer edges are colored blue imply that two angles incident to N and two angles incident to S are bicolored. Hence, $\Lambda = 4n + 4$.

Moreover, as T has a quadrangular outer face and triangular inner faces, Euler's relation ensures that T has $2n + 2$ inner faces. For each inner face, two cases can arise: either the three sides have the same color, or two sides are of one color and one side is of the other color. In the first (second) case, the face has 0 (2, respectively) bicolored angles. As there are $2n + 2$ inner faces and $\Lambda = 4n + 4$, the pigeonhole principle implies that all inner faces have a contribution of 2 to the number of bicolored angles, which concludes the proof. $\square$

Lemma 1 ensures that Ψ is a mapping from $\mathcal{E}$ to $\mathcal{O}$. In addition, it is clear that Ψ is injective. To prove Proposition 2, it remains to prove that Φ and Ψ are surjective. As Φ is injective, it is sufficient to show that $\Psi \circ \Phi$ is surjective. Thus, given $O \in \mathcal{O}$, we have to find $Z \in \mathcal{B}$ such that $\Psi \circ \Phi(Z) = O$.

Computing the preimage of an α_0-orientation. We now describe a method to compute a transversal pair of bipolar orientations Z consistent with a given α_0-orientation O, i.e., such that $\Psi \circ \Phi(Z) = O$. The algorithm makes use of a sweeping process to orient and color the inner edges of T; a simple (i.e., not self-intersecting) path $\mathcal{P}$ of inner edges of T going from W to E is maintained, the path moving progressively toward the vertex S (at the end, the path is $\mathcal{P} = W \to S \to E$). We require that the following invariants are satisfied throughout the sweeping process.

(1) Each vertex of $\mathcal{P}\backslash\{W, E\}$ has two outgoing edges on each side of $\mathcal{P}$ for the α_0-orientation O.

(2) The inner edges of T already oriented and colored are those on the left of $\mathcal{P}$.

(3) Condition C1' holds around each inner vertex of T on the left of $\mathcal{P}$.

(4) A partial version of C1' holds around each vertex v of $\mathcal{P}\backslash\{W, E\}$. The edges incident to v on the left of $\mathcal{P}$ form in clockwise order: a possibly empty interval of ingoing blue edges, a non-empty interval of outgoing red edges, and a possibly empty interval of outgoing blue edges.

(5) All edges already oriented and colored and incident to N, E, S, W are ingoing red, ingoing blue, outgoing red, and outgoing blue, respectively.

(6) The edges of T already colored (and oriented) are consistent with the α_0-orientation O, i.e., for each angle (v, f) delimited by two edges of T already colored, the corresponding edge of $Q(T)$ is going out of v iff the angle is bicolored.

At first we need a technical result ensuring that Invariant (1) is sufficient for a path to be simple.

Lemma 2 *Let T be an irreducible triangulation and $Q(T)$ be the angular graph of T, endowed with an α_0-orientation O. Let $\mathcal{P}$ be a path of inner edges of T from W to E such that each vertex of $\mathcal{P}\backslash\{W, E\}$ has outdegree 2 on each side of $\mathcal{P}$ for the α_0-orientation. Then the path $\mathcal{P}$ is simple.*

PROOF. Assume that the path $\mathcal{P}$ loops into a circuit; and consider an inclusion-minimal such circuit $\mathcal{C} = (v_0, v_1, \ldots, v_k = v_0)$. We define $n_\bullet$, $n_\circ$ and e as the numbers of black vertices (i.e., vertices of T), white vertices and edges of $Q(T)$ inside $\mathcal{C}$. As T is triangulated and $\mathcal{C}$ has length k, Euler's relation ensures that T has $2n_\bullet + k - 2$ faces inside $\mathcal{C}$, i.e., $n_\circ = 2n_\bullet + k - 2$. Counting the edges of $Q(T)$ inside $\mathcal{C}$ according to their incident white vertex gives (i): $e = 3n_\circ = 6n_\bullet + 3k - 6$. The edges of $Q(T)$ inside $\mathcal{C}$ can also be counted according to their origin for the α_0-orientation. Each vertex of $\mathcal{C}$ —except possibly the self-intersection vertex v_0— has outdegree 2 in the interior of $\mathcal{C}$ for the α_0-orientation. Hence, (ii): $e = 4n_\bullet + n_\circ + 2k - 2 + \delta = 6n_\bullet + 3k - 4 + \delta$, where $\delta \geq 0$ is the outdegree of v_0 inside $\mathcal{C}$. Taking $(ii) - (i)$ yields $\delta = -2$, a contradiction. $\square$

The path $\mathcal{P}$ is initialized with all neighbours of N, from W to E. In addition, all inner edges incident to N are initially colored red and directed toward N, see Figure 5(b). The invariants (1)-to-(6) are clearly true at the initial step.

Let us introduce some terminology in order to describe the sweeping process. Thoughout the process, the vertices of $\mathcal{P}$ are ordered from left to right, with
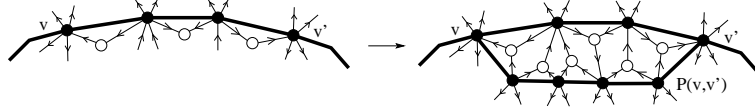
Fig. 3. An admissible pair v, v' of vertices, and the matching path $P(v, v')$.

W as leftmost and E as rightmost vertex. Given v, v' a pair of vertices on $\mathcal{P}$ —with v on the left of v'— the part of $\mathcal{P}$ going from v to v' is denoted by $[v, v']$. For each vertex w on $\mathcal{P}$, let $f_1, \ldots, f_k$ be the sequence of faces of T incident to w on the right of $\mathcal{P}$, taken in counterclockwise order. The edge of $Q(T)$ associated to the angle (w, f_1) (angle (w, f_k)) is denoted by $\epsilon_{\text{left}}(w)$ (by $\epsilon_{\text{right}}(w)$, respectively). A pair of vertices v, v' on $\mathcal{P}$ —with v on the left of v'— is called *admissible* if $\epsilon_{\text{right}}(v)$ is ingoing at v, $\epsilon_{\text{left}}(v')$ is ingoing at v', and for each vertex $w \in [v, v'] \backslash \{v, v'\}$, the edges $\epsilon_{\text{left}}(w)$ and $\epsilon_{\text{right}}(w)$ are going out of w. Clearly an admissible pair always exists: take a pair v, v' of vertices on $\mathcal{P}$ that satisfy $\{\epsilon_{\text{right}}(v)$ ingoing at v, $\epsilon_{\text{left}}(v')$ ingoing at $v'\}$ and are closest possible. Notice that two vertices v, v' forming an admissible pair are not neighbours on $\mathcal{P}$ (otherwise the white vertex associated to the face on the right of $[v, v']$ would have outdegree > 1). Let $w_0 = v, w_1, \ldots, w_k, w_{k+1} = v'$ ($k \geq 1$) be the sequence of vertices of $[v, v']$. The *matching path* of v, v' is the path of edges of T that starts at v, visits the neighbours of $w_1, w_2, \ldots, w_k$ on the right of $\mathcal{P}$, and finishes at v'. The matching path of v, v' is denoted by $P(v, v')$. Let $\mathcal{P}'$ be the path obtained from $\mathcal{P}$ when substituting $[v, v']$ by $P(v, v')$. As shown in Figure 3, the path $\mathcal{P}'$ goes from W to E and each vertex of $\mathcal{P}' \backslash \{W, E\}$ has outdegree 2 on each side of $\mathcal{P}'$ for the α_0-orientation O. Hence the path $\mathcal{P}'$ is simple according to Lemma 2. Moreover, by definition of $P(v, v')$, all edges of T in the region enclosed by $[v, v']$ and $P(v, v')$ connect a vertex of $[v, v'] \backslash \{v, v'\}$ to a vertex of $P(v, v') \backslash \{v, v'\}$, see Figure 4.

We can now describe the operations performed at each step of the sweeping process, as shown in Figure 4.

- Choose an admissible pair v, v' of vertices on $\mathcal{P}$.
- Color blue and orient from left to right all edges of $[v, v']$.
- Color red all edges inside the area enclosed by $[v, v']$ and $P(v, v')$, and orient these edges from $P(v, v')$ to $[v, v']$.
- Update the path $\mathcal{P}$, the part $[v, v']$ being replaced by $P(v, v')$.

According to the discussion above, the path $\mathcal{P}$ is still simple after these operations, and satisfies Invariant (1). All the other invariants (2)-to-(6) are easily shown to remain satisfied, as illustrated in Figure 4. At the end, the path $\mathcal{P}$ is equal to $W \to S \to E$. The invariants (3) and (5) ensure that the directions and colors of the inner edges of T form a transversal pair of bipolar orientations Z; and Invariant (6) ensures that Z is consistent with the α_0-orientation O, i.e., $\Psi \circ \Phi(Z) = O$, see also Figure 5 for a complete execution of the algorithm. This concludes the proof of Proposition 2.
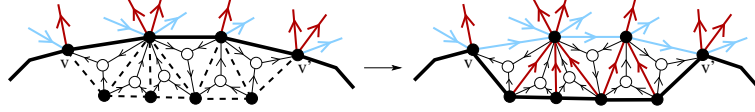
11

Fig. 4. The update step of the iterative algorithm to find the preimage of an α_0-orientation.
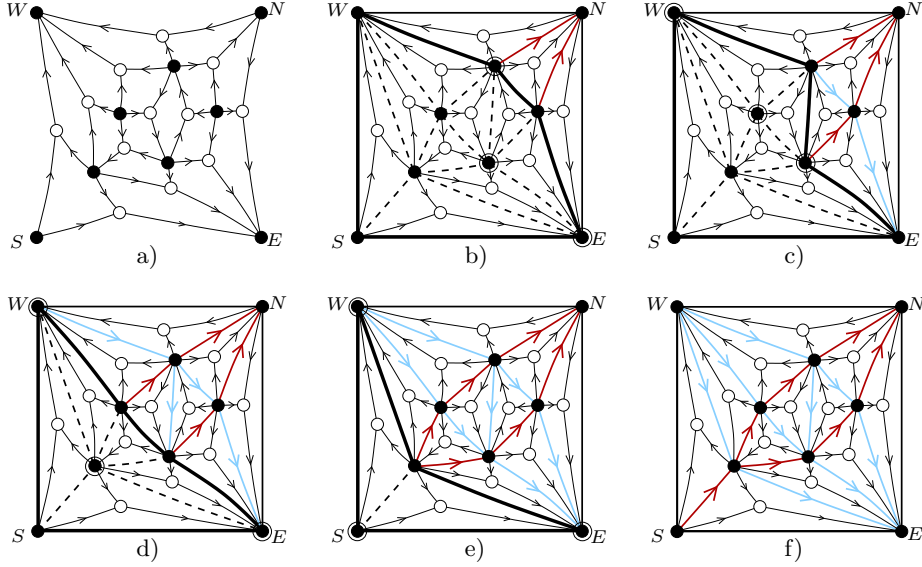

Fig. 5. The complete execution of the algorithm calculating the preimage of an α_0-orientation O. At each step, the vertices of the matching path $P(v, v')$ are surrounded.

3.3 Essential circuits of an α_0-orientation

Proposition 2 ensures that the set $\mathcal{E}$ of transversal edge-partitions of an irreducible triangulation is a distributive lattice, as $\mathcal{E}$ is in bijection with the distributive lattice formed by the α_0-orientations of the angular graph. By definition, the flip operation on $\mathcal{E}$ is the effect of a flip operation on $\mathcal{O}$ via the bijection. Recall that a flip operation on an α-orientation consists in reversing a clockwise essential circuit (circuit with no chordal path). Hence, to describe the flip operation on $\mathcal{E}$, we have to characterise the essential circuits of an α_0-orientation. For this purpose, we introduce the concept of *straight path*.

Consider an irreducible triangulation T endowed with a transversal pair of bipolar orientations. Color blue the four outer edges of T and orient them from W to E. The conditions C1' and C2' ensure that there are four possible types for a bicolored angle (e, e') of T, with e' following e in cw order: (outgoing red, outgoing blue) or (outgoing blue, ingoing red) or (ingoing red, ingoing blue) or (ingoing blue, outgoing red). The type of an edge of $Q(T)$ corresponding to a bicolored angle of T (i.e., an edge going out of a black vertex) is defined as the type of the bicolored angle. For such an edge e, the *straight path* of e

is the oriented path $\mathcal{P}$ of edges of $Q(T)$ that starts at e and such that each edge of $\mathcal{P}$ going out of a black vertex has the same type as e. Such a path is unique, as there is a unique choice for the outgoing edge at a white vertex (each white vertex has outdegree 1).

Lemma 3 *The straight path $\mathcal{P}$ of an edge $e \in Q(T)$ going out of a black vertex is simple and ends at an outer black vertex of $Q(T)$.*

PROOF. Notice that the conditions of a transversal pair of bipolar orientations remain satisfied if the directions of the edges of one color are reversed and then the colors of all inner edges are switched; hence the edge e can be assumed to have type (outgoing red, outgoing blue) without loss of generality. Let $v_0, v_1, v_2, \ldots, v_i, \ldots$ be the sequence of vertices of the straight path $\mathcal{P}$ of e, so that the even indices correspond to black vertices of $Q(T)$ and the odd indices correspond to white vertices of $Q(T)$. Observe that, for $k \geq 0$, the black vertices v_{2k} and v_{2k+2} are adjacent in T and the edge (v_{2k}, v_{2k+2}) is either outgoing red or outgoing blue. Hence $(v_0, v_2, v_4, \ldots, v_{2k}, \ldots)$ is an oriented path of T, so that it is simple, according to Proposition 1. Hence, $\mathcal{P}$ does not pass twice by the same black vertex, i.e., $v_{2k} \neq v_{2k+2}$ for $k \neq k'$; and $\mathcal{P}$ neither passes twice by the same white vertex (otherwise $v_{2k+1} = v_{2k+3}$ for $k \neq k'$ would imply $v_{2k+2} = v_{2k+4}$ by unicity of the outgoing edge at each white vertex, a contradiction). Thus $\mathcal{P}$ is a simple path, so that it ends at a black vertex of $Q(T)$ having no outgoing edge of type (outgoing red, outgoing blue), i.e., $\mathcal{P}$ ends at an outer black vertex of $Q(T)$. $\square$

Proposition 3 *Given T an irreducible triangulation and $Q(T)$ its angular graph endowed with an α_0-orientation X, an essential clockwise circuit $\mathcal{C}$ of X satisfies either of the two following configurations,*

- *The circuit $\mathcal{C}$ is the boundary of a (quadrangular) inner face of $Q(T)$, see Figure 6(a).*
- *The circuit $\mathcal{C}$ has length 8. The four black vertices of $\mathcal{C}$ have no outgoing edge inside $\mathcal{C}$. The four white vertices of $\mathcal{C}$ have their unique incident edge not on $\mathcal{C}$ inside $\mathcal{C}$, see Figure 6(b) for an example.*

PROOF. First we claim that no edge of $Q(T)$ inside $\mathcal{C}$ has its origin on $\mathcal{C}$; indeed the straight path construction ensures that such an edge could be extended to a chordal path of $\mathcal{C}$, which is impossible. We define $n_\bullet$, $n_\circ$ and e as the number of black vertices, white vertices and edges inside $\mathcal{C}$. We denote by $2k$ the number of vertices on $\mathcal{C}$, so that there are k black and k white vertices on $\mathcal{C}$. Euler's relation and the fact that all inner faces of $Q(T)$ are quadrangular ensure that (i) : $e = 2(n_\bullet + n_\circ) + k - 2$. As each white vertex of $Q(T)$ has degree 3, a white vertex on $\mathcal{C}$ has a unique incident edge not on $\mathcal{C}$. Let l be the
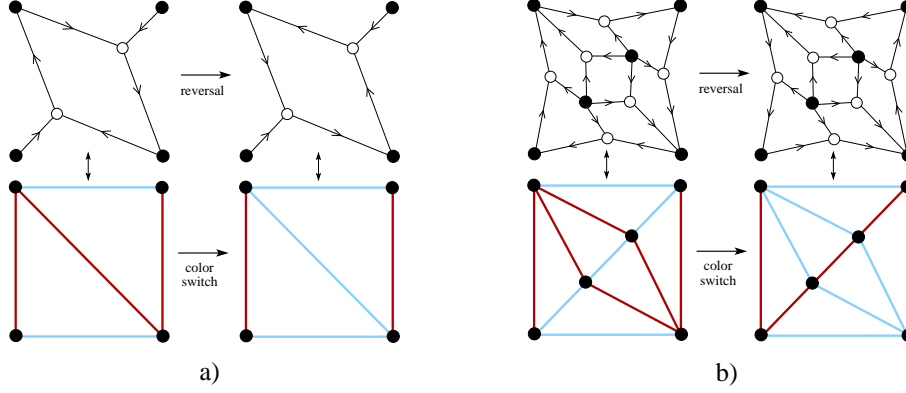
Fig. 6. The two possible configurations of an essential clockwise circuit $\mathcal{C}$ of $Q(T)$. In each case, an alternating 4-cycle is associated to the circuit; and reversing the circuit orientation corresponds to switching the edge colors inside the alternating 4-cycle.

number of white vertices such that this incident edge is inside $\mathcal{C}$ (notice that $l \leq k$). Counting the edges inside $\mathcal{C}$ according to their incident white vertex gives $(ii):\ e = 3n_\circ + l$. Edges inside $\mathcal{C}$ can also be counted according to their origin for the α_0-orientation. As no edge inside $\mathcal{C}$ has its origin on $\mathcal{C}$, we have $(iii):\ e = 4n_\bullet + n_\circ$. Taking $2(i) - (ii) - (iii)$ yields $l = 2k - 4$. As k is a positive integer and l is a nonnegative integer satisfying $l \leq k$, the only possible values for l and k are $\{k = 4, l = 4\}$, $\{k = 3, l = 2\}$, and $\{k = 2, l = 0\}$. It is easily seen that the case $\{k = 3, l = 2\}$ would correspond to a separating 3-cycle. Hence, the only possible cases are $\{k = 2, l = 0\}$ and $\{k = 4, l = 4\}$, shown in Figure 6(a) and 6(b), respectively. The first case corresponds to a circuit of length 4, which has to be the boundary of a face, as the angular graph of an irreducible triangulation (more generally, of a 3-connected plane graph) is well known to have no filled 4-cycle. □

3.4   Flip operation on transversal structures

As we prove now, the essential circuits of α_0-orientations correspond to specific patterns on transversal structures, making it possible to have a simple geometric interpretation of the flip operation, formulated directly on the transversal structure.

Given T an irreducible triangulation endowed with a transversal edge-partition, we define an *alternating 4-cycle* as a cycle $\mathcal{C} = (e_1, e_2, e_3, e_4)$ of 4 edges of T that are color-alternating (i.e., two adjacent edges of $\mathcal{C}$ have different colors). Given a vertex v on $\mathcal{C}$, we call left-edge (right-edge) of v the edge of $\mathcal{C}$ starting from v and having the exterior of $\mathcal{C}$ on its left (on its right, respectively).

Lemma 4 *An alternating 4-cycle $\mathcal{C}$ in a transversal structure satisfies either of the two following configurations.*

- *All edges inside $\mathcal{C}$ and incident to a vertex v of $\mathcal{C}$ have the color of the left-edge of v. Then $\mathcal{C}$ is called a* left alternating 4-cycle
- *All edges inside $\mathcal{C}$ and incident to a vertex v of $\mathcal{C}$ have the color of the right-edge of v. Then $\mathcal{C}$ is called a* right alternating 4-cycle.

PROOF. Let k be the number of vertices inside $\mathcal{C}$. Condition C1 ensures that there are $4k$ bicolored angles incident to a vertex inside $\mathcal{C}$. Moreover, Euler's relation ensures that there are $2k + 2$ faces inside $\mathcal{C}$. Hence Lemma 1 implies that there are $4k + 4$ bicolored angles inside $\mathcal{C}$. As a consequence, there are four bicolored angles inside $\mathcal{C}$ that are incident to a vertex of $\mathcal{C}$. As $\mathcal{C}$ is alternating, each of the four vertices of $\mathcal{C}$ is incident to at least one bicolored angle. Hence the pigeonhole principle implies that each vertex of $\mathcal{C}$ is incident to one bicolored angle inside $\mathcal{C}$. Moreover, each of the four inner faces (f_1, f_2, f_3, f_4) inside $\mathcal{C}$ and incident to an edge of $\mathcal{C}$ has two bicolored angles, according to Lemma 1. As such a face f_i has two angles incident to vertices of $\mathcal{C}$, at least one bicolored angle of f_i is incident to a vertex of $\mathcal{C}$. As there are four bicolored angles inside $\mathcal{C}$ incident to vertices of $\mathcal{C}$, the pigeonhole principle ensures that each face f_1, f_2, f_3, f_4 has exactly one bicolored angle incident to a vertex of $\mathcal{C}$. As each vertex of $\mathcal{C}$ is incident to one bicolored angle inside $\mathcal{C}$, these angles are in the same direction. If they start from $\mathcal{C}$ in clockwise (counterclockwise) direction, then $\mathcal{C}$ is a right alternating 4-cycle (left alternating 4-cycle, respectively). $\qquad\square$

Theorem 2 *The set of transversal edge-partitions of a fixed irreducible triangulation is a distributive lattice. The flip operation consists in switching the edge-colors inside a right alternating 4-cycle, turning it into a left alternating 4-cycle.*

PROOF. The two possible configurations for an essential clockwise circuit of an α_0-orientation are represented respectively in Figure 6(a) and Figure 6(b). In each case, the essential circuit of the angular graph corresponds to a right alternating 4-cycle on T. Notice that a clockwise face of $Q(T)$ corresponds to the vertex-empty alternating 4-cycle, whereas essential circuits of length 8 correspond to all possible alternating 4-cycles with at least one vertex in their interior (Figure 6(a) gives an example). As shown in Figure 6, the effect of reversing a clockwise essential circuit of the angular graph is clearly a color switch of the edges inside the associated alternating 4-cycle. $\qquad\square$

Definition. Given an irreducible triangulation T, the transversal structure of
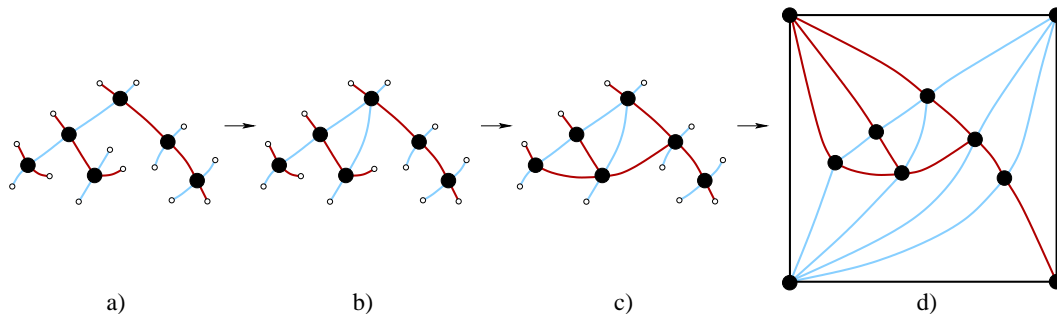
Fig. 7. The execution of the closure mapping on an example.

T with no right alternating 4-cycle is called *minimal*, as it is at the bottom of the distributive lattice.

4   Bijection with ternary trees

This section focuses on the description of a bijection between ternary trees and irreducible dissections, where the minimal transversal structure (for the distributive lattice) plays a crucial role. The mapping from ternary trees to irreducible triangulations relies on so-called *closure operations*, as introduced by G. Schaeffer in his PhD [25].

4.1   The closure mapping: from trees to triangulations

A *plane tree* is a plane graph with a unique face (the outer face). A ternary tree is a plane tree with vertex degrees in $\{1, 4\}$. Vertices of degree 4 are called *nodes* and vertices of degree 1 are called *leaves*. An edge of a ternary tree is called a *closed edge* if it connects two nodes and is called a *stem* if it connects a node and a leaf. It will be convenient to consider closed edges as made of two opposite half-edges meeting at the middle of the edge, whereas stems will be considered as made of a unique half-edge incident to the node and having not (yet) an opposite half-edge. A ternary tree is *rooted* by marking one leaf. The root allows us to distinguish the four neighbours of each node, taken in ccw-order, into a parent (the neighbour in the direction of the root), a left-child, a middle-child, and a right-child. Thus, our definition of rooted ternary trees corresponds to the classical definition, where each node has three ordered children.

Starting from a ternary tree, the three steps to construct an irreducible triangulation are: local closure, partial closure and complete closure. Perform a counterclockwise walk alongside a ternary tree A (imagine an ant walking around A with the infinite face on its right). If a stem s and then two sides of

16

closed edges e_1 and e_2 are successively encountered during the traversal, create a half-edge opposite to the stem s and incident to the farthest extremity of e_2, so as to *close* a triangular face. This operation is called a *local closure*, see the transition Figure 7(a)-(b).

The figure obtained in this way differs from A by the presence of a triangular face and, more importantly, a stem s of A has become a closed edge, i.e., an edge made of two half-edges. Each time a sequence (stem, closed edge, closed edge) is found in ccw around the outer face of the current figure, we perform a local closure, update the figure, and restart, until no local closure is possible. This greedy execution of local closures is called the *partial closure* of A, see Figure 7(c). It is easily shown that the figure F obtained by performing the partial closure of A does not depend on the order of execution of the local closures. Indeed, a cyclic parenthesis word is associated to the counter-clockwise walk alongside the tree, with an opening parenthesis of weight 2 for a stem and a closing parenthesis for a side of closed edge; the future local closures correspond to matchings of the parenthesis word.

At the end of the partial closure, the number n_s of unmatched stems and the number n_e of sides of closed edges incident to the outer face of F satisfy the relation $n_s - n_e = 4$. Indeed, this relation is satisfied on A because a ternary tree with n nodes has $n - 1$ closed edges and $2n + 2$ leaves (as proved by induction on the number of nodes); and the relation $n_s - n_e = 4$ remains satisfied throughout the partial closure, as each local closure decreases n_s and n_e by 1. When no local closure is possible anymore, two consecutive unmatched stems on the boundary of the outer face of F are separated by at most one closed edge. Hence, the relation $n_s = n_e + 4$ implies that the unmatched stems of F are partitioned into four intervals I_1, I_2, I_3, I_4, where two consecutive stems of an interval are separated by one closed edge, and the last stem of I_i is incident to the same vertex as the first stem of $I_{(i+1) \bmod 4}$, see Figure 7(c).

The last step of the closure mapping, called *complete closure*, consists of the following operations. Draw a 4-gon (v_1, v_2, v_3, v_4) outside of F; for $i \in \{1, 2, 3, 4\}$, create an opposite half-edge for each stem s of the interval I_i, the new half-edge being incident to the vertex v_i. Clearly, this process creates only triangular faces, so that it yields a triangulation of the 4-gon, see Figure 7(d).

Let us now explain how the closure mapping is related to transversal edge-partitions. A ternary tree A is said to be *edge-bicolored* if each edge of A (closed edge or stem) is given a color —red or blue— such that any angle incident to a node of A is bicolored, see Figure 7(a). Such a bicoloration, which is unique up to the choice of the colors, is called the *edge-bicoloration* of A.

Lemma 5 *Let A be a ternary tree endowed with its edge-bicoloration. The following invariant is maintained throughout the partial closure of A,*

17

(I): any angle incident to the outer face of the current figure is bicolored.

PROOF. By definition of the edge-bicoloration, (I) is true on A. We claim that (I) remains satisfied after each local closure. Indeed, let (s, e_1, e_2) be the succession (stem, closed edge, closed edge) intervening in the local closure, let v be the extremity of e_2 farthest from s, and let e_3 be the ccw follower of e_2 around v. Invariant (I) implies that s and e_2 have the same color. As we give to the new created half-edge h the same color as its opposite half-edge s (in order to have unicolored edges), h and e_2 have the same color. The effect of the local closure on the angles of the outer face is the following: the angle (e_1, e_2) disappears from the outer face, and the angle (e_2, e_3) is replaced by the angle (h, e_3). As e_2 has the same color as h, the bicolored angle (e_2, e_3) is replaced by the bicolored angle (h, e_3), so that (I) remains true after the local closure. $\qquad\square$

It also follows from this proof that Condition C1 remains satisfied throughout the partial closure, because the number of bicolored angles around each node is not increased, and is initially equal to 4. At the end of the partial closure, Invariant (I) ensures that all stems of the intervals I_1 and I_3 are of one color, and all stems of the intervals I_2 and I_4 are of the other color. Hence, Condition C2 is satisfied after the complete closure, see Figure 7(d). Thus, the closure maps the edge-bicoloration of A to a transversal edge-partition of the obtained triangulation of the 4-gon, which is in fact the minimal one:

Proposition 4 *The closure of a ternary tree A with n nodes is an irreducible triangulation T with n inner vertices. The closure maps the edge-bicoloration of A to the minimal transversal edge-partition of T.*

PROOF. Assume that T has a separating 3-cycle $\mathcal{C}$. Observe that Lemma 1 was stated and proved without the irreducibility condition. Hence, when the four outer edges of T are colored blue, each inner face of T has exactly two bicolored angles. Let $k \geq 1$ be the number of vertices inside $\mathcal{C}$. Euler's relation implies that the interior of $\mathcal{C}$ contains $2k + 1$ faces, so that there are $4k + 2$ bicolored angles inside $\mathcal{C}$, according to Lemma 1. Moreover, Condition C1 implies that there are 4k bicolored angles incident to a vertex that is in the interior of $\mathcal{C}$. Hence there are exactly two bicolored angles inside $\mathcal{C}$ incident to a vertex of $\mathcal{C}$. However, for each of the three edges $\{e_1, e_2, e_3\}$ of $\mathcal{C}$, the face incident to e_i in the interior of $\mathcal{C}$ has at least one of its two bicolored angles incident to e_i. Hence, there are at least three bicolored angles inside $\mathcal{C}$ and incident to a vertex of $\mathcal{C}$, a contradiction.

Now we show that the transversal edge-partition of T induced by the closure mapping is minimal, i.e., has no right alternating 4-cycle. Let $\mathcal{C}$ be an alter-
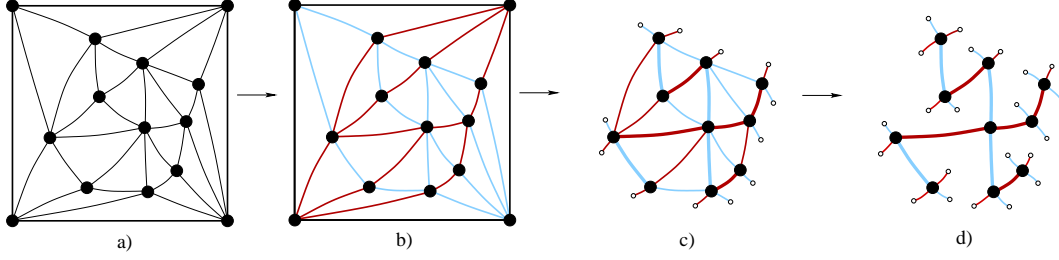
Fig. 8. The opening algorithm performed on an example.

nating 4-cycle of T. This cycle has been closed during a local closure involving one of the four edges of $\mathcal{C}$. Let e be this edge and let v be the origin of the stem whose completion has created the edge e. The fact that the closure of a stem is always performed with the infinite face on its right ensures that e is the right-edge of v on $\mathcal{C}$, as defined in Section 3.4. A second observation following from Invariant (I) is that, when a stem s is merged, the angle formed by s and by the edge following s in counterclockwise order around the origin of s is a bicolored angle. This ensures that $\mathcal{C}$ is a left alternating 4-cycle. $\qquad\square$

4.2 Inverse mapping: the opening

In this section, we describe the inverse of the closure mapping, from irreducible triangulations to ternary trees. As we have seen in the proof of Lemma 5, during a local closure, the newly created half-edge h has the same color as the clockwise-consecutive half-edge around the origin of h. Hence, for each half-edge h incident to an inner vertex of T,

- if the angle formed by h and its cw-consecutive half-edge is unicolored, then h has been created during a local closure,
- if the angle is bicolored, then h is one of the 4 original half-edges of A incident to v.

This property indicates how to inverse the closure mapping. Given an irreducible triangulation T, the *opening* of T consists of the following steps, illustrated in Figure 8.

(1) Endow T with its minimal transversal edge-partition.
(2) Remove the outer quadrangle of T and all half-edges of T incident to a vertex of the quadrangle.
(3) Remove all the half-edges whose clockwise-consecutive half-edge has the same color.

The following lemma is a direct consequence of the definition of the opening mapping:

19

Lemma 6 *Let A be a ternary tree and let T be the irreducible triangulation obtained by performing the closure of A. Then the opening of T is A.*

Hence, the closure Φ and the opening Ψ are such that $\Psi \circ \Phi = \mathrm{Id}$. To prove that the opening and the closure mapping are mutually inverse, it remains to prove that $\Phi \circ \Psi = \mathrm{Id}$, which is more difficult and is done in two steps:
1) show that the opening of an irreducible triangulation T is a ternary tree,
2) show that the closure of this ternary tree is T.

The first step is to define an orientation of the half-edges of T induced by the minimal transversal edge-partition. Each half-edge of T is associated to the angle on its right (looking from the incident vertex). We orient the half-edges of T toward (outward of) their incident vertex if the associated angle is unicolored (bicolored, respectively); with the restriction that half-edges on the outer quadrangle are leaved unoriented and all angles incident to an outer vertex are considered as unicolored. This yields an orientation of the half-edges of T, which is called the *4-orientation* of T. Each inner vertex of T is incident to four bicolored angle, hence has outdegree 4 in the 4-orientation. By definition of the 4-orientation, the opening of an irreducible triangulation consists in removing the outer 4-gon and all ingoing half-edges.

Let e be an inner edge of T. An important remark is that the two half-edges of e can not be simultaneously directed toward their respective incident vertex (otherwise, the 4-cycle $\mathcal{C}$ bordering the two triangular faces incident to e would be a right alternating 4-cycle, a contradiction).

Hence, only two cases arise for an inner edge e of T.

- If one half-edge of e is ingoing, then e is called a *stem-edge*. A stem-edge can be considered as simply oriented (both half-edges have the same direction) for the 4-orientation.
- If the two half-edges of e are outgoing, then e is called a *tree-edge*. A tree-edge can be considered as a bi-oriented edge for the 4-orientation.

We define a clockwise circuit of the 4-orientation of T as a simple cycle $\mathcal{C}$ of inner edges of T such that each edge of $\mathcal{C}$ is either a tree-edge (i.e., a bi-oriented edge) or a stem-edge having the interior of $\mathcal{C}$ on its right.

Lemma 7 *The 4-orientation of T has no clockwise circuit.*

PROOF. Assume there exists a clockwise circuit $\mathcal{C}$ in the 4-orientation of T. For a vertex v on $\mathcal{C}$, we denote by h_v the half-edge of $\mathcal{C}$ going out of v when doing a clockwise traversal of $\mathcal{C}$, and by e_v the edge of $Q(T)$ following h_v in clockwise order around v (notice that e_v is the most counterclockwise edge of $Q(T)$ incident to v inside $\mathcal{C}$). As $\mathcal{C}$ is a clockwise circuit for the 4-orientation
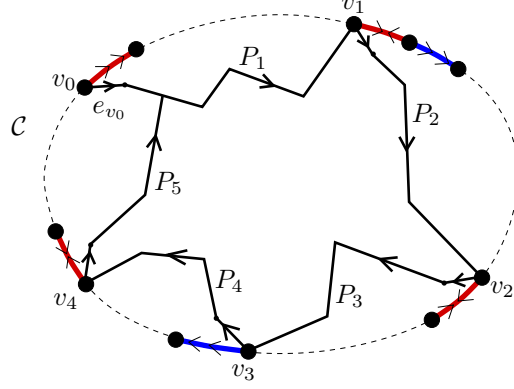
Fig. 9. The existence of a clockwise circuit in the 4-orientation of T implies the existence of a clockwise circuit in the minimal α_0-orientation of $Q(T)$.

of T, h_v is directed outward of v. Hence the angle θ on the right of h_v is a bicolored angle for the minimal transversal edge-partition of T, so that e_v is going out of v for the minimal α_0-orientation $O_{\min}$ of $Q(T)$.

We use this observation to build iteratively a clockwise circuit of $O_{\min}$ (see Figure 9), yielding a contradiction. Let v_0 be a vertex on $\mathcal{C}$ and $\mathcal{P}(v_0)$ the straight path starting at e_{v_0}, as defined in Section 3.3, for the minimal α_0-orientation $O_{\min}$ of $Q(T)$. Lemma 3 ensures that $\mathcal{P}(v_0)$ is simple and ends at an outer vertex of $Q(T)$. In particular, $\mathcal{P}(v_0)$ has to reach $\mathcal{C}$ at a vertex v_1 different from v_0. We denote by P_1 the part of $\mathcal{P}(v_0)$ between v_0 and v_1, by Λ_1 the part of $\mathcal{C}$ between v_1 and v_0, and by $\mathcal{C}_1$ the cycle obtained by concatenating P_1 and Λ_1. Let $\mathcal{P}(v_1)$ be the straight path starting at e_{v_1}. The fact that e_{v_1} is the most counterclockwise incident edge of v_1 in the interior of $\mathcal{C}$ ensures that $\mathcal{P}(v_1)$ starts in the interior of $\mathcal{C}_1$. The path $\mathcal{P}(v_1)$ has to reach $\mathcal{C}_1$ at a vertex $v_2 \neq v_1$. We denote by P_2 the part of $\mathcal{P}(v_1)$ between v_1 and v_2. If v_2 belongs to P_1, then the concatenation of the part of P_1 between v_2 and v_1 and of P_2 is a clockwise circuit of $O_{\min}$, a contradiction. Hence v_2 is on Λ_1 strictly between v_1 and v_0. We denote by $\overline{P}_2$ the concatenation of P_1 and P_2, and by Λ_2 the part of $\mathcal{C}$ going from v_2 to v_0. As v_2 is strictly between v_1 and v_0, Λ_2 is strictly included in Λ_1. We denote by $\mathcal{C}_2$ the cycle made of the concatenation of $\overline{P}_2$ and Λ_2. Similarly as for the path $\mathcal{P}(v_1)$, the path $\mathcal{P}(v_2)$ must start in the interior of $\mathcal{C}_2$.

Then we continue iteratively, see Figure 9. At each step k, we consider the straight path $\mathcal{P}(v_k)$ starting at e_{v_k}. This path starts in the interior of the cycle $\mathcal{C}_k$, and reaches again $\mathcal{C}_k$ at a vertex v_{k+1}. The vertex v_{k+1} can not belong to $\overline{P}_k$, otherwise a clockwise circuit of $O_{\min}$ would be created. Hence v_{k+1} is strictly between v_k and v_0 on $\mathcal{C}$, i.e., is in $\Lambda_k \backslash \{v_k, v_0\}$. In particular the path Λ_{k+1} going from v_{k+1} to v_0 on $\mathcal{C}$, is strictly included in the path Λ_k going from v_k to v_0 on $\mathcal{C}$. Thus, Λ_k shrinks strictly at each step. Hence, there must be a step k_0 where $\mathcal{P}(v_{k_0})$ reaches $\mathcal{C}_{k_0}$ at a vertex on $\overline{P}_{k_0}$, thus creating a clockwise circuit of $O_{\min}$, a contradiction. $\qquad \square$

Lemma 8 *The tree-edges of T form a tree spanning the inner vertices of T.*

PROOF. Denote by H the graph consisting of the tree-edges of T and their incident vertices. A first observation is that H has no cycle, as such a cycle of bi-oriented edges of T would be a clockwise circuit in the 4-orientation of T. Let n be the number of inner vertices of T. Observe that H can not be incident to the outer vertices of T, so that H can cover at most the set of inner vertices of T. A well-known result of graph theory ensures that an acyclic graph H having $n - 1$ edges and covering a subset of an n-vertex set V is a tree covering exactly all vertices of V. Hence it remains to show that H has $(n - 1)$ edges. Let s be the number of stem-edges and t be the number of tree-edges of T. As T has n inner vertices, there are $4n$ outgoing half-edges in the 4-orientation of T. Moreover, each stem-edge has contribution 1 and each tree-edge has contribution 2 to the number of outgoing half-edges. Hence, $s + 2t = 4n$. Finally, Euler's relation ensures that T has $(3n + 1)$ inner edges, so that $s + t = 3n + 1$. These two equalities ensure that $t = n - 1$, which concludes the proof that H is a tree spanning the inner vertices. □

Lemma 9 *The opening of an irreducible triangulation T is a ternary tree.*

PROOF. As we have seen from the definition of the 4-orientation, the opening of an irreducible triangulation consists in removing the outer 4-gon and all ingoing half-edges. The figure obtained in this way consists of the tree-edges, which form a spanning tree according to Lemma 8, and of the half-edges that have lost their opposite half-edge. The edges of the first and second type correspond respectively to the closed edges and to the stems of the tree. In addition, after removing all ingoing half-edges, each vertex has degree 4, so that the tree satisfies the degree-conditions of a ternary tree. □

Lemma 10 *Let T be an irreducible triangulation and let A be the ternary tree obtained by doing the opening of T. Then the closure of A is T.*

PROOF. First it is clear that the complete closure (transition between Figure 7(c) and Figure 7(d)) is the inverse of Step 2 of the opening algorithm. Let F be the figure obtained from T after Step 2 of the opening mapping.

To prove that the partial closure of A is F, it is sufficient to find a chronological order of deletion of the ingoing half-edges of F (for the 4-orientation) such that the inverse of each half-edge deletion is a local closure. A local closure satisfies the property that the new created half-edge h has the outer face on its right when h is traversed toward its incident vertex. Thus the half-edge h chosen to be disconnected at step k must be incident to the outer face of the current

figure F_k, with F_k on the right when h is traversed toward its incident vertex. We claim that there always exists such a half-edge as long as there remain stem-edges in F_k. In that case, F_k contains the spanning tree H made of the tree-edges of T plus at least one stem-edge. Hence F_k contains at least a cycle, thus a simple cycle $\mathcal{C}$ can be extracted from the boundary of F_k. There is at least one stem-edge e on $\mathcal{C}$, because no cycle is formed by tree-edges only. We claim that e has the outer face of F_k on its right, so that the ingoing half-edge h of e is a candidate to be deleted (indeed if e had the outer face of F_k on its left, the concatenation of e and of the path of tree-edges connecting the two extremities of e would be a clockwise circuit in the 4-orientation, a contradiction). As discussed above, the inverse operation of the deletion of h is a local closure, which concludes the proof. □

Finally, Lemma 6 and Lemma 10 yield the following theorem:

Theorem 3 (bijection) *For $n \geq 1$, the closure mapping is a bijection between the set of ternary trees with n nodes and the set of irreducible triangulations with n inner vertices. The inverse mapping of the closure is the opening.*

Theorem 4 (bijection, rooted version) *For $n \geq 1$, the closure mapping induces a $(2n+2)$-to-4 correspondence between the set $\mathcal{A}'_n$ of rooted ternary trees with n nodes and the set $\mathcal{T}'_n$ of rooted irreducible triangulations with n inner vertices. In other words, $\mathcal{A}'_n \times \{1, \ldots, 4\}$ is in bijection with $\mathcal{T}'_n \times \{1, \ldots, 2n+2\}$.*

PROOF. It can easily be proved by induction on the number of nodes that a ternary tree with n nodes has $2n+2$ leaves. Hence, when rooting the ternary tree obtained by doing the opening of a triangulation in $\mathcal{T}'_n$, there are $2n+2$ possibilities to place the root. Conversely, starting from a rooted ternary tree with n nodes, there are four possibilities to place the root on the irreducible triangulation obtained by doing the closure of the tree, because the root has to be placed on one of the four outer edges. □

4.3 Application to counting triangulations

In this section we focus on the application to exact enumeration. The key point is that the bijection reduces the task of counting irreducible triangulations to the much easier task of counting ternary trees. Then, we show that the enumeration of irreducible triangulations naturally leads to the enumeration of rooted 4-connected triangulations, which are closely related; the ingredients are generating functions and a decomposition of a rooted irreducible triangulation as a sequence of rooted 4-connected triangulations.

4.3.1 Counting irreducible triangulations

As a first direct application, the bijection with ternary trees yields counting formulas for irreducible triangulations.

Proposition 5 (counting irreducible triangulations) *For $n \geq 1$, the number of rooted irreducible triangulations with n inner vertices is*

$$|\mathcal{T}_n'| = 4\frac{(3n)!}{n!(2n+2)!}.$$

The number of unrooted irreducible triangulations with n inner vertices is

$$|\mathcal{T}_n| = \frac{(3n)!}{n!(2n+2)!} + \frac{1}{2}\frac{(3k)!}{k!(2k+1)!} \qquad \text{if } n \equiv 0 \mod 2 \ \ [n = 2k],$$

$$|\mathcal{T}_n| = \frac{(3n)!}{n!(2n+2)!} + \frac{1}{2}\frac{(3k+1)!}{k!(2k+2)!} + \frac{1}{2}\frac{(3k')!}{k'!(2k'+1)!} \qquad \text{if } n \equiv 1 \mod 4$$

$$[n = 2k+1 = 4k'+1],$$

$$|\mathcal{T}_n| = \frac{(3n)!}{n!(2n+2)!} + \frac{1}{2}\frac{(3k+1)!}{k!(2k+2)!} \qquad \text{if } n \equiv 3 \mod 4 \ \ [n = 1+2k].$$

PROOF. The enumerative formula follows from $|\mathcal{T}_n'| = \frac{4}{2n+2}|\mathcal{A}_n'|$ and from the well known fact that $|\mathcal{A}_n'| = (3n)!/((2n+1)!n!)$, which can be derived from the Lagrange inversion formula applied to the generating function $A(z) = z(1 + A(z))^3$. The formula for $|\mathcal{T}_n| = |\mathcal{A}_n|$ follows from the enumeration of unrooted ternary trees, which is easily obtained by considering the possible rotation symmetries (order 2 around a vertex or an edge, order 4 around a vertex). □

The formula for *rooted* irreducible triangulations can easily be obtained from the series counting rooted triangulations of the 4-gon by using a composition scheme, see [30]. To our knowledge, the counting formula for *unrooted* irreducible triangulations is new. However, a composition scheme should make it possible to count irreducible triangulations with a given rotation symmetry (order 2 around a vertex or an edge, order 4 around a vertex), starting from triangulations of the 4-gon with a given rotation symmetry, which have been counted by Brown [7].

4.3.2 Counting rooted 4-connected triangulations

A graph is said to be *4-connected* if it has more than three vertices and if at least four vertices have to be removed to disconnect it. In this section we de-

rive the enumeration of rooted 4-connected triangulations from the counting formula for rooted irreducible triangulations. The idea is to translate a decomposition linking these two families of rooted triangulations to an equation linking their generating functions. The net result we obtain is an explicit formula for the generating function of rooted 4-connected triangulations (Proposition 6). We take advantage of the well known property that a triangulation is 4-connected iff the interior of any 3-cycle, except for the outer triangle, is a face. In all this section, we denote by $\mathcal{T}' = \cup_n \mathcal{T}'_n$ and by $\mathcal{C}' = \cup_n \mathcal{C}'_n$ the sets of rooted irreducible triangulations and of rooted 4-connected triangulations counted with respect to the number of inner vertices.

Observe the close connection between the definition of 4-connected triangulations and irreducible triangulations. In particular, for $n \geq 2$, the operation of removing the root edge of an object of $\mathcal{C}'_n$ and carrying the root on the counterclockwise-consecutive edge is an injective mapping from $\mathcal{C}'_n$ to $\mathcal{T}'_{n-1}$. However, given $T \in \mathcal{T}'_{n-1}$, the inverse edge-adding operation can create a separating 3-cycle if there exists an internal path of length 2 connecting the origin of the root of T to the vertex diametrically opposed in the outer (quadrangular) face of T. Objects of $\mathcal{T}'$ having no such internal path are said to be *undecomposable* and their set, counted with respect to the number n of inner vertices, is denoted by $\mathcal{U}' = \cup_n \mathcal{U}'_n$. The above discussion ensures that $\mathcal{C}'_n$ is in bijection with $\mathcal{U}'_{n-1}$ for $n \geq 2$. In addition, a maximal decomposition of an object $\gamma \in \mathcal{T}'$ along the above mentioned interior paths of length 2 ensures that γ is a sequence of objects of $\mathcal{U}'$. Precisely, the graph enclosed by two consecutive paths of length 2 is either an undecomposable triangulation or is a quadrangle with a unique interior edge that connects the middles of the two paths. This leads to the equation

$$T(z) + 1 = \frac{U(z) + 1}{1 - z(U(z) + 1)}, \tag{1}$$

where $T(z) = \sum |\mathcal{T}'_n| z^n$ and $U(z) = \sum |\mathcal{U}'_n| z^n$ are respectively the series counting $\mathcal{T}'$ and $\mathcal{U}'$ with respect to the number of inner vertices.

Proposition 6 *The series $C(z)$ counting rooted 4-connected triangulations by their number of inner vertices has the following expression,*

$$C(z) = \frac{z(A(z) - A(z)^2 + 1)}{1 + z(A(z) - A(z)^2 + 1)}, \tag{2}$$

where $A(z) = z(1 + A(z))^3$ is the series counting rooted ternary trees by their number of nodes.

PROOF. As $\mathcal{C}'_n$ is in bijection with $\mathcal{U}'_{n-1}$ for $n \geq 2$ and as the unique 4-connected triangulation with less than 5 vertices is the tetrahedron, we have

$C(z) = z(U(z) + 1)$. Hence, Equation (1) yields $C(z) = z(T(z) + 1)/(1 + z(T(z) + 1))$. Thus, it remains to provide an expression of $T(z)$ in terms of the series $A(z) = \sum A_n z^n$ counting rooted ternary trees by their number of nodes. We define respectively the sets $\overline{\mathcal{A}_n}$ and $\widehat{\mathcal{A}_n}$ of ternary trees with n nodes and having the following marks: a closed edge is marked and oriented for objects of $\overline{\mathcal{A}_n}$, a closed edge is marked and oriented and a leaf is marked for objects of $\widehat{\mathcal{A}_n}$. As a ternary tree with n nodes has $n - 1$ closed edges and $2n + 2$ leaves, we have $A_n \cdot 2(n - 1) = |\widehat{\mathcal{A}_n}| = |\overline{\mathcal{A}_n}|(2n + 2)$, so that $|\overline{\mathcal{A}_n}| = 2\frac{n-1}{2n+2} A_n$. In addition, given a ternary tree with a marked oriented edge, the operation of cutting the marked edge produces an ordered pair of rooted ternary trees. Hence, the series counting $\overline{\mathcal{A}} := \cup_n \overline{\mathcal{A}_n}$ with respect to the number of nodes is $A(z)^2$. Finally, we know from Corollary 4 that $T'_n = \frac{4}{2n+2} A_n$. Hence, we also have $|\mathcal{T}'_n| = \frac{2n+2}{2n+2} A_n - 2\frac{n-1}{2n+2} A_n = A_n - |\overline{\mathcal{A}_n}|$, from which we conclude that $T(z) = A(z) - A(z)^2$. Finally, to obtain the expression of $C(z)$ in terms of $A(z)$, substitute $T(z)$ by $A(z) - A(z)^2$ in the expression $C(z) = z(T(z) + 1)/(1 + z(T(z) + 1))$. $\qquad\square$

From Expression (2), the coefficients of $C(z)$ can be quickly extracted: $C(z) = z + z^3 + 3\,z^4 + 12\,z^5 + 52\,z^6 + 241\,z^7 + \mathcal{O}(z^8)$. The first coefficients match with the values given by the formula $c_n = \dfrac{1}{n}\sum_{i=1}^{n}(-1)^i \dbinom{3n - i - 1}{n - i}\dbinom{2i}{2}$ for the number of rooted 4-connected triangulations with $n - 1$ inner vertices. This formula was found by Tutte [30] using more complicated algebraic methods.

5 Application to straight-line drawing

A *straight-line drawing* of a plane graph G is a planar embedding of G such that vertices are on an integer grid $[0, W] \times [0, H]$ and the edges are represented as segments. The integers W and H are called the width and the height of the grid. As described next, transversal structures give rise to a simple straight-line drawing algorithm, based on face-counting operations.

5.1 Face-counting algorithm

Given an irreducible triangulation T endowed with a transversal pair of bipolar orientations, we define the *red-map* of T as the plane graph T_r obtained from T by removing all blue edges, see Figure 1(e). The four outer edges of T_r are additionally oriented from S to N, so that T_r is endowed with a bipolar orientation, with S and N as poles (according to Corollary 1). Similarly, define the *blue-map* T_b of T as the map obtained by removing the red edges of T.

The four outer edges of T_b are additionally oriented from W to E, so that T_b is endowed with a bipolar orientation. Given a vertex v of T, we define the *leftmost outgoing red path* of v as the oriented path $\mathcal{P}_r^{\mathrm{out}}(v)$ of edges of the red-map starting at v and such that each edge of the path is the leftmost outgoing edge at its origin. As the orientation of the red-map is bipolar, $\mathcal{P}_r^{\mathrm{out}}(v)$ is simple and ends at N. We also define the *rightmost ingoing red path* of v as the path $\mathcal{P}_r^{\mathrm{in}}(v)$ starting from v and such that each edge of $\mathcal{P}_r^{\mathrm{in}}(v)$ is the rightmost ingoing edge at its end-vertex. This path is also acyclic and ends at S. The *separating red path* of v, denoted by $\mathcal{P}_r(v)$, is the concatenation of $\mathcal{P}_r^{\mathrm{out}}(v)$ and $\mathcal{P}_r^{\mathrm{in}}(v)$. The oriented (simple) path $\mathcal{P}_r(v)$ goes from S to N, thus inducing a bipartition of the inner faces of T_r, whether a face is on the left or on the right of $\mathcal{P}_r(v)$. Similarly, we define the *leftmost outgoing blue path* $\mathcal{P}_b^{\mathrm{out}}(v)$, the *rightmost ingoing blue path* $\mathcal{P}_b^{\mathrm{in}}(v)$, and write $\mathcal{P}_b(v)$ for the concatenation of $\mathcal{P}_b^{\mathrm{out}}(v)$ and $\mathcal{P}_b^{\mathrm{in}}(v)$, called the *separating blue path* of v.

The following facts easily follow from the definition of the paths.

Fact 1 *Let v and v' be two different vertices of T.*

- *The paths $\mathcal{P}_r^{\mathrm{out}}(v)$ and $\mathcal{P}_r^{\mathrm{out}}(v')$ do not cross each other, they join at a vertex v'' and then are equal between v'' and N.*
- *The paths $\mathcal{P}_r^{\mathrm{in}}(v)$ and $\mathcal{P}_r^{\mathrm{in}}(v')$ do not cross each other, they join at a vertex v'' and then are equal between v'' and S.*
- *The paths $\mathcal{P}_r^{\mathrm{out}}(v)$ and $\mathcal{P}_r^{\mathrm{in}}(v')$ do not cross each other. In addition, $\mathcal{P}_r^{\mathrm{in}}(v')$ can not join $\mathcal{P}_r^{\mathrm{out}}(v)$ from the left of $\mathcal{P}_r^{\mathrm{out}}(v)$.*

These facts easily imply that the separating red paths of two distinct vertices do not cross each other, and the same holds for separating blue paths.

The face-counting algorithm TRANSVERSALDRAW runs as follows —see Figure 10—, the input being the combinatorial description of an irreducible triangulation endowed with a transversal structure:

TRANSVERSALDRAW(T):
- For each vertex v of T,
  - the abscissa of v is the number of inner faces of the red-map T_r on the left of $\mathcal{P}_r(v)$,
  - the ordinate of v is the number of inner faces of the blue-map T_b on the right of $\mathcal{P}_b(v)$.
- Place the vertices of T on an integer grid accordingly, and link adjacent vertices by segments.

Theorem 5 *Given T an irreducible triangulation endowed with a transversal structure, TRANSVERSALDRAW(T) outputs a straight-line drawing of T. The half-perimeter of the grid is $n - 1$ if T has n vertices. The embedded edges*
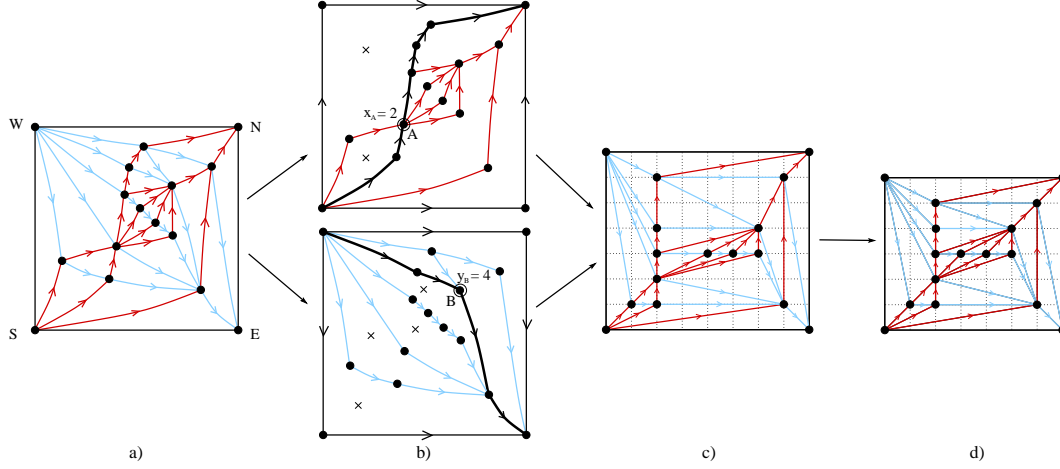
Fig. 10. The execution of TRANSVERSALDRAW (Fig.a-to-c) on an example. The deletion of unused coordinates —COMPACTTRANSVERSALDRAW— is shown in Fig.d.

satisfy the following orientation property*:*

- *red edges are oriented from bottom to top and weakly oriented from left to right,*
- *blue edges are oriented from left to right and weakly oriented from top to bottom.*

The algorithm TRANSVERSALDRAW *can be implemented to run in linear time.*

PROOF. Let us first prove the orientation property. Let $e = (v, v')$ be a red edge of T, directed from v to v'. As discussed above, $\mathcal{P}_b(v)$ and $\mathcal{P}_b(v')$ do not cross each other. Clearly Condition C1' implies that $\mathcal{P}_b(v)$ is on the right of $\mathcal{P}_b(v')$. Hence, the ordinate of v' is greater than the ordinate of v. Then, proving that the abscissa of v' is at least as large as the abscissa of v reduces to proving that $\mathcal{P}_r(v)$ is not on the right of $\mathcal{P}_r(v')$. This follows from two easy observations: $\mathcal{P}_r^{\text{out}}(v)$ is (weakly) on the left of $\mathcal{P}_r^{\text{out}}(v')$; and $\mathcal{P}_r^{\text{in}}(v')$ is (weakly) on the right of $\mathcal{P}_r^{\text{in}}(v')$. Similarly, the blue edges are oriented from left to right and weakly oriented from top to bottom.

The proof that the embedding is a straight-line drawing relies on the orientation property and on the fact that the red and the blue edges are combinatorially transversal. To carry out the proof, we use a sweeping process akin to the iterative algorithm used to find the preimage of an α_0-orientation in Section 3.2. The idea consists in maintaining an oriented blue path $\mathcal{P}$ from W to E called the *sweeping path*, such that the following invariant is maintained, (I_{draw}): "the embedding of the edges of T that are not topologically on the right of $\mathcal{P}$ is a straight-line drawing delimited to the top by the embedding of (W, N), to the right by the embedding of (N, E), and to the bottom-left by the embedding of $\mathcal{P}$". The sweeping path is initially (W, N, E) and then
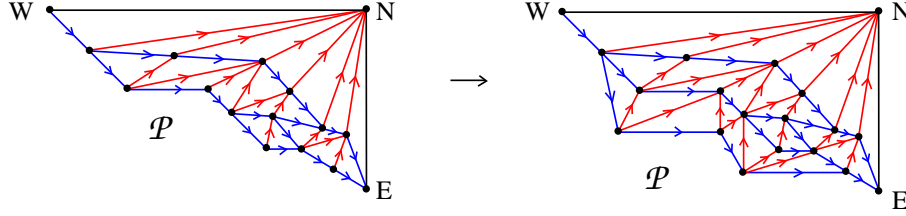
28

Fig. 11. The direction properties of the edges ensure that the drawing is planar, step after step.

"moves" toward S. At each step, an inner face f of the blue-map of T is chosen, such that the left lateral path of f is included in $\mathcal{P}$. Then, $\mathcal{P}$ is updated by replacing the left lateral path of f by the right lateral path of f. Thus $\mathcal{P}$ remains an oriented path from W to E and is moved toward the bottom-left corner of the embedding, i.e., the vertex S. The fact that the invariants are maintained is easily checked from the orientation property and the fact that all the red edges inside f connect transversally the two lateral blue paths, see Figure 11. At the end, the sweeping path is equal to (W, S, E), so that the invariant (I_{draw}) exactly implies that the embedding of T is planar.

We show now that the half-perimeter of the grid is equal to $n - 1$ if T has n vertices. By definition of TRANSVERSALDRAW, the minimal abscissa is 0 and the maximal abscissa is equal to the number of inner faces f_r of the red-map. Similarly, the minimal ordinate is 0 and the maximal ordinate is equal to the number of inner faces f_b of the blue-map. Hence, the half-perimeter is equal to $f_r + f_b$. We write respectively e_r and e_b for the number of edges of T_r and T_b. Euler's relation ensures that the total number e of edges of T is $3n - 7$. Hence, $e_b + e_r = e + 4 = 3n - 3$. In addition, Euler's relation, applied respectively to T_r and T_b, ensures that $n + f_r = e_r + 1$ and $n + f_b = e_b + 1$. Thus, $f_r + f_b = e_r + e_b - 2n + 2 = n - 1$.

Finally, a linear implementation is obtained by suitably performing the face-counting operations. For each inner vertex v of T, consider the rightmost outgoing red path, the leftmost outgoing red path, the leftmost ingoing red path, and the rightmost ingoing red path. These four paths P_1, P_2, P_3, P_4 partition the set of inner faces of the red-map T_r into four areas $\mathcal{U}(v)$, $\mathcal{L}(v)$, $\mathcal{D}(v)$, and $\mathcal{S}(v)$, that are respectively enclosed by (P_1, P_2), (P_2, P_3), (P_3, P_4), and (P_4, P_1). Let $U(v)$, $L(v)$, $D(v)$ and $S(v)$ be the numbers of inner faces of T_r in each of the four areas. The quantities $U(v)$ are easily computed in one pass, by doing a traversal of the vertices of T from S to N. Similarly, the quantities $D(v)$ are computed in one pass doing a traversal from N to S. Then, the quantities $L(v)$ are computed in one pass (using $D(v)$ and $U(v)$) by doing a traversal from W to E. Finally, the abscissas of all vertices can be computed using $\mathrm{Abs}(v) = D(v) + L(v)$. □

5.2 Compaction step by coordinate-deletions

Consider an irreducible triangulation T endowed with a transversal structure and embedded using the face-counting algorithm TRANSVERSALDRAW. As observed in Figure 10(c), some coordinate-lines might be unoccupied. Hence, a natural further step is to delete the unused coordinates, yielding a more compact drawing, as illustrated in Figure 10(d). The corresponding algorithm is called COMPACTTRANSVERSALDRAW.

Theorem 6 *Given T an irreducible triangulation endowed with a transversal structure, COMPACTTRANSVERSALDRAW(T) outputs a straight-line drawing of T. The half-perimeter of the grid is at most $n-1$ if T has n vertices. The embedded edges satisfy the same orientation property as TRANSVERSALDRAW(T), and the algorithm can be implemented to run in linear time.*

PROOF. Clearly, the orientation property of the red and of the blue edges remains satisfied after coordinate-deletions. Notice that the planarity of TRANSVERSALDRAW has been proved using only the orientation property, so that the same proof of planarity works as well for COMPACTTRANSVERSALDRAW. The linear complexity follows from the easy property that coordinate-deletions can be performed in linear time. □

5.3 Drawing with the minimal transversal structure

Notice that, in the definition of both TRANSVERSALDRAW and COMPACT-TRANSVERSALDRAW, the irreducible triangulation is already equipped with a transversal structure. A complete straight-line drawing algorithm for irreducible triangulations is thus obtained by computing a transversal structure first, and then launching the face-counting algorithm TRANSVERSALDRAW, possibly followed by the additional step of deletion of unused coordinates (COMPACTTRANSVERSALDRAW). A natural choice for the transversal structure computed is to take the minimal one for the distributive lattice, with the further advantage of making it possible to perform an average-case analysis of the grid size, as we will see in Section 6. The minimal transversal structure can be computed in linear time: 1) compute a transversal structure in linear time using an algorithm by Kant and He [20], 2) make the transversal structure minimal by iterated circuit reversals on the associated α_0-orientations of the angular graph (the fact that the overall complexity of circuit reversions is linear easily follows from ideas presented in [21]). Alternatively, a linear algorithm COMPUTEMINIMAL computing directly the minimal transversal structure is described in the PhD of the author [16]. We define the following straight-line drawing algorithms for irreducible triangulations.
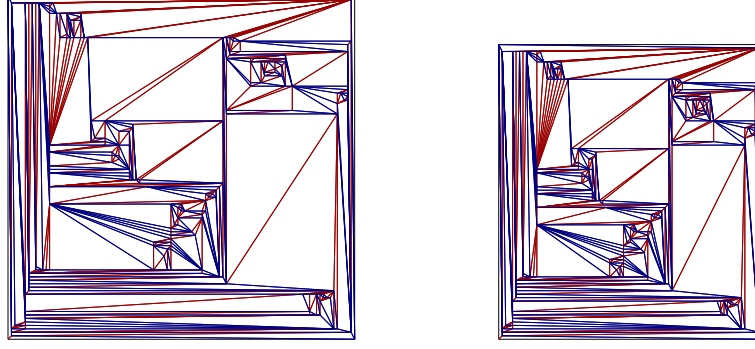
Fig. 12. A random triangulation with 200 vertices embedded with the algorithms DRAW and COMPACTDRAW.

DRAW(T):
- Call COMPUTEMINIMAL(T)
- Call TRANSVERSALDRAW(T)

COMPACTDRAW(T):
- Call COMPUTEMINIMAL(T)
- Call COMPACTTRANSVERSALDRAW(T)

Simulations on random irreducible triangulations of large size ($n \approx 50000$) indicate that the grid size is always approximately $\frac{n}{2} \times \frac{n}{2}$ with DRAW and $\frac{n}{2}(1-\alpha) \times \frac{n}{2}(1-\alpha)$ with COMPACTTRANSVERSALDRAW, for some constant $\alpha \approx 0.18$, see Figure 12 for an example. It turns out that the size of the grid can be analyzed thanks to the closure mapping presented in Section 4. In the same way, as described by Bonichon *et al* [5], the bijection of [24] makes it possible to perform a probabilistic analysis of the grid size for straight-line drawing algorithms based on Schnyder woods [4].

Finally let us mention that a more general procedure of compaction is described in the PhD of the author [16]: (1) delete a set $\mathcal{E}$ of red and blue edges such that the red acyclic orientation and the blue acyclic orientation remain bipolar, (2) call the face-counting algorithm TRANSVERSALDRAW, (3) re-insert the edges of $\mathcal{E}$. It is shown in [16] that the drawing thus obtained is planar and has semi-perimeter $n-1-|\mathcal{E}|$. (However the orientation properties might not be satisfied.)

6   Analysis of the grid size of the drawing algorithms

In this section, we explain how the closure mapping gives rise to a precise probabilistic analysis of the grid size of the algorithms DRAW and COMPACT-DRAW. The main result is as follows.

Theorem 7 *Let T be a rooted irreducible triangulation with n inner vertices taken uniformly at random. Let $W \times H$ be the grid size of* DRAW*(T) and let $W_c \times H_c$ be the grid size of* COMPACTDRAW*(T). Then, the following results*

hold asymptotically with high probability, up to fluctuations of order $\sqrt{n}$,

$$W \times H \simeq \frac{n}{2} \times \frac{n}{2}, \qquad W_c \times H_c \simeq \frac{11}{27}n \times \frac{11}{27}n.$$

Let us mention that the concentration property also holds in a so-called ϵ-*formulation*: "for any $\epsilon > 0$, the probability that W or H are outside of $[(1/2 - \epsilon)n, (1/2 + \epsilon)n]$ and the probability that W_c or H_c are outside of $[(11/27 - \epsilon)n, (11/27 + \epsilon)n]$ are asymptotically exponentially small".

We first deal with the analysis of W and H, i.e., the width and the height of $\text{Draw}(T)$. This task is rather easy and allows us to introduce some tools that will also be used to analyze the grid size of $\text{CompactDraw}(T)$, namely generating functions and the so-called quasi power theorem. In the sequel, the set of rooted irreducible triangulations with n inner vertices is denoted by $\mathcal{T}_n'$.

6.1   Analysis of the grid size of $\text{Draw}(T)$

Given $T \in \mathcal{T}_n'$ endowed with its minimal transversal pair of bipolar orientations, the width of the grid of $\text{Draw}(T)$ is, by definition, the number of inner faces of the red-map T_r. Let e_r and f_r be the numbers of inner edges and of inner faces of T_r. Euler's relation applied to T_r ensures that $e_r = n + f_r - 1$. By definition of the opening, e_r is equal to the number of red edges (including the stems) of the edge-bicolored ternary tree obtained by doing the opening of T. An easy adaptation of the proof of Theorem 4 ensures that the uniform distribution on $\mathcal{T}_n'$ is mapped by the opening to the uniform distribution on rooted edge-bicolored ternary tree with n nodes. These observations lead to the following statement:

Fact 2 *The distribution of the width of $\text{Draw}(T)$ for T uniformly sampled in $\mathcal{T}_n'$ is equal to the distribution of $e_r - n - 1$, where e_r is the number of red edges of a uniformly sampled rooted edge-bicolored ternary tree with n nodes.*

We denote by $\mathcal{R}$ (*resp.* $\mathcal{B}$) the set of rooted edge-bicolored ternary trees whose root leaf is incident to a red stem (*resp.* blue stem). For a rooted edge-bicolored ternary tree γ, we denote by $|\gamma|$ the number of nodes of γ and by $\xi(\gamma)$ the number of red edges of γ. We define the generating functions

$$R(z, u) = \sum_{\gamma \in \mathcal{R}} z^{|\gamma|} u^{\xi(\gamma)}, \quad B(z, u) = \sum_{\gamma \in \mathcal{B}} z^{|\gamma|} u^{\xi(\gamma)}$$

that count the set $\mathcal{R}$ and the set $\mathcal{B}$ with respect to the number of nodes and the number of red edges. The generating function $E(z, u)$ counting rooted edge-bicolored ternary trees with respect to the number of nodes and the number

32

of red edges is thus equal to $R(z, u) + B(z, u)$. The classical decomposition of a rooted ternary tree at the root node into three subtrees translates to the following equation system:

$$\begin{cases} R(z, u) = zu\,(1 + B(z, u))\,(u + R(z, u))\,(1 + B(z, u)), \\ B(z, u) = z\,(u + R(z, u))\,(1 + B(z, u))\,(u + R(z, u)). \end{cases}$$

This system is polynomial in R, B, z and u, so that one derives, by algebraic elimination, a trivariate polynomial $P(E, z, u)$ that satisfies $P(E(z, u), z, u) = 0$, i.e., $E(z, u)$ is an algebraic series.

We state now an adaptation of the quasi-power theorem [14, Thm. IX.7, Cor. IX.1], [11] for algebraic series. This theorem makes it possible to analyse the distribution of the number of red edges in a rooted edge-bicolored ternary tree.

Theorem 8 (Algebraic quasi power theorem) *Let $E(z, u)$ be the generating function of a combinatorial class $\mathcal{E}$, where the variable z marks the size $|\gamma|$ of an object $\gamma \in \mathcal{E}$ and the variable u marks a parameter ξ, i.e., $E(z, u) = \sum_{\gamma \in \mathcal{E}} z^{|\gamma|} u^{\xi(\gamma)}$. Assume that $E(z, u)$ is an algebraic series, i.e., there exists a trivariate polynomial $P(E, z, u)$ with rational coefficients such that $P(E(z, u), z, u) = 0$. Consider the polynomial system:*

$$\text{(S)} := \{P(E, z, u) = 0, \quad \frac{\partial P}{\partial E}(E, z, u) = 0\}.$$

Assume that (S) has a solution $\{\tau, \rho\}$ at $u = 1$ such that τ and ρ are positive real values and there exists no other complex solution (E, z) of (S) at $u = 1$ such that $|z| \leq \rho$. Assume further that the derivative-condition $\{\frac{\partial^2 P}{\partial E^2}(\tau, \rho, 1) \neq 0, \frac{\partial P}{\partial z}(\tau, \rho, 1) \neq 0\}$ is satisfied. Assume also that the Jacobian of (S) with respect to E and z does not vanish at $(\tau, \rho, 1)$, i.e.,

$$\det \begin{pmatrix} \frac{\partial P}{\partial E}(\tau, \rho, 1) & \frac{\partial^2 P}{\partial E^2}(\tau, \rho, 1) \\ \frac{\partial P}{\partial z}(\tau, \rho, 1) & \frac{\partial^2 P}{\partial z \partial E}(\tau, \rho, 1) \end{pmatrix} \neq 0.$$

Then there exists a unique pair of algebraic series $(\tau(u), \rho(u))$ such that

$$\{P(\tau(u), \rho(u), u) = 0, \frac{\partial P}{\partial E}(\tau(u), \rho(u), u) = 0, \tau(1) = \tau, \rho(1) = \rho\}.$$

Finally, assume that $\rho(u)$ satisfies the mean condition $\rho'(1) \neq 0$ and variance condition $\rho''(1)\rho(1) + \rho'(1)\rho(1) - \rho'(1)^2 \neq 0$. Then, for γ taken uniformly at random in the set of objects of $\mathcal{E}$ of size n, the random variable $X_n = \xi(\gamma)$ is asymptotically equal to μn, up to Gaussian fluctuations of order $\sigma \sqrt{n}$, where

$$\mu = -\frac{\rho'(1)}{\rho(1)}, \quad \sigma^2 = -\frac{\rho''(1)}{\rho(1)} - \frac{\rho'(1)}{\rho(1)} + \left(\frac{\rho'(1)}{\rho(1)}\right)^2.$$

In other words, $(X_n - \mu n)/(\sigma\sqrt{n})$ converges to the normal law.

This theorem, despite a rather long statement and several conditions to check, is easy to apply in practice. In our case, $E(z, 1)$ is clearly equal to $2A(z)$ with $A(z) = z(1 + A(z))^3$ the generating function of rooted ternary trees. Hence, at $u = 1$, we have $P(E, z) = E/2 - z(1 + E/2)^3$. Using a computer algebra software, the solution of (S) at $u = 1$ is found to be $\{\tau = 1, \rho = 4/27\}$. The derivative condition and the Jacobian condition are then easily checked. The algebraic function $\rho(u)$ is obtained by taking the resultant of the two equations of (S) with respect to E; then the factor $Q(z, u)$ of the resultant such that $Q(\rho, 1) = 0$ gives an algebraic equation for $\rho(u)$, i.e., $Q(\rho(u), u) = 0$. From the algebraic equation $Q(\rho(u), u) = 0$, the derivative and second derivative of $\rho(u)$ at $u = 1$ are readily calculated. To calculate $\rho'(1)$, we derive $Q(\rho(u), u) = 0$ and find $\frac{\partial Q}{\partial z}(\rho(u), u)\rho'(u) + \frac{\partial Q}{\partial u}(\rho(u), u) = 0$, so that $\rho'(1) = -\frac{\partial Q}{\partial u}(\rho, \tau)/\frac{\partial Q}{\partial z}(\rho, \tau)$. The mean condition and variance condition are readily checked and we find $\mu = 3/2$. Hence, the number e_r of red edges in a random rooted edge-bicolored ternary tree is asymptotically equal to $3n/2$, up to fluctuations of order $\sqrt{n}$. This result, together with Fact 2, ensure that the width of $\text{DRAW}(T)$, for T taken uniformly at random in $\mathcal{T}'_n$, is asymptotically equal to $n/2$ up to fluctuations of order $\sqrt{n}$. The result is the same for the height H of $\text{DRAW}(T)$, by stability of $\mathcal{T}'_n$ and of the algorithm DRAW under the $\pi/2$-clockwise rotation. Thus, we have proved the statement of Theorem 7 on the distribution of the grid size of $\text{DRAW}(T)$.

6.2 Analysis of the grid size of COMPACTDRAW(T)

6.2.1 Overview of the proof

Now we concentrate on the distribution of the grid width of COMPACT-DRAW(T), for T taken uniformly at random in $\mathcal{T}'_n$ (the analysis of the grid height is similar). By definition, the width of COMPACTDRAW(T) is $W - \Delta_{\text{abs}(T)}$, where W is the width of $\text{DRAW}(T)$ and Δ_{abs} is the number of abscissas not used by $\text{DRAW}(T)$. We have already proved that W is asymptotically equal to $n/2$ up to fluctuations of order $\sqrt{n}$. Hence, to obtain the statement of Theorem 7 about W_c, it is sufficient to prove that $\Delta_{\text{abs}(T)}$ is asymptotically equal to $5n/54$ up to fluctuations of order $\sqrt{n}$. The steps of the proof are the following: first, the abscissas that are not used by $\text{DRAW}(T)$ are associated with specific edges of T; then these edges of T are shown to correspond to specific red edges of the ternary tree obtained by doing the opening of T. Finally, using generating functions and the algebraic quasi power theorem, it is proved that the number of such edges in a random rooted edge-bicolored ternary tree with n nodes is asymptotically equal to $5n/54$, up to fluctuations of order $\sqrt{n}$.

6.2.2 Characterisation of unused abscissas as particular edges of T

We analyze the number of unused abscissas of $\text{DRAW}(T)$, i.e., the number of vertical lines of the grid that bear no vertex of T. Recall that the width of the grid of $\text{DRAW}(T)$ is the number of inner faces f_r of the red-map T_r of T, where T_r is obtained by computing the minimal transversal pair of bipolar orientations of T and then removing all blue edges. By definition of $\text{TRANSVERSAL-DRAW}(T)$, the abscissa $\text{Abs}(v)$ of a vertex v of T is obtained by associating with v an oriented path $\mathcal{P}_r(v)$ of the red-map, called the separating red path of v (as defined in Section 5), and then counting the number of inner faces of T_r on the left of $\mathcal{P}_r(v)$. We show first that each abscissa-candidate $1 \leq i \leq f_r$ can be associated with a unique inner face of T_r —for which we write $\text{Abs}(f) = i$— such that the existence of a vertex v of T with $\text{Abs}(v) = i$ only depends on the configuration of the edge at the bottom-right corner of f.

Let us start with a few definitions. Given $e = (v, v')$ an edge of T_r directed from v to v', the *separating red path* of e, denoted by $\mathcal{P}_r(e)$, is the path obtained by concatenating the leftmost outgoing red path of v', the edge e, and the rightmost ingoing red path of v. We recall that, for each inner face f of T_r, there are two vertices s_f and t_f of f such that the boundary of f consists of a left and a right lateral paths going from s_f to t_f. The *separating red path* of f, denoted by $\mathcal{P}_r(f)$, is defined as the concatenation of the leftmost outgoing red path of t_f, of the right lateral path $P_{\text{right}}(f)$ of f, and of the rightmost ingoing red path of s_f. Observe that the first edge e_f of $P_{\text{right}}(f)$, called bottom-right edge of f, satisfies $\mathcal{P}_r(f) = \mathcal{P}_r(e_f)$.

Lemma 11 *Let e and e' be two different edges of the red-map T_r. Then the separating paths $\mathcal{P}_r(e)$ and $\mathcal{P}_r(e')$ do not cross each other.*

PROOF. First it is easily checked that the bipolar orientation on red edges (and similarly the one on blue edges) has no transitive edge; otherwise, there would be a face of the red-map with one lateral path reduced to an edge, in contradiction with the presence of transversal blue edges within the face. Hence only three cases can arise: either e' and $\mathcal{P}_r(e)$ do not intersect, or they intersect at a unique extremity of e', or e' is on $\mathcal{P}_r(e)$. Fact 1 allows us to check that $\mathcal{P}_r(e)$ and $\mathcal{P}_r(e')$ do not cross each other for each of the three cases. □

Recall that the separating red path of a face f is the separating red path of its bottom-right edge. Thus, Lemma 11 ensures that the separating red paths $\mathcal{P}_r(f)$ and $\mathcal{P}_r(f')$ of two different inner faces f and f' of T_r do not cross each other. In addition, it is easy to see that they are different using the fact that the bottom-right edge of a face is not the leftmost outgoing red edge at its origin. As a consequence, $\text{Abs}(f) \neq \text{Abs}(f')$. There are f_r inner faces in T_r,

each inner face f clearly satisfying $1 \leq \mathrm{Abs}(f) \leq f_r$. Hence the pigeonhole principle yields:

Fact 3 *For $1 \leq i \leq f_r$, there exists a unique inner face f of T such that $\mathrm{Abs}(f) = i$.*

Thus, the inner faces of T_r are strictly ordered from left to right according to their associated abscissa.

Lemma 12 *The separating red paths of an edge e of the red-map T_r and of a vertex $v \in T$ do not cross each other. In addition, given e an edge of T_r, there exists a vertex v such that $\mathcal{P}_r(v) = \mathcal{P}_r(e)$ iff either e is the rightmost ingoing red edge at its end-vertex or e is the leftmost outgoing red edge at its origin.*

PROOF. The fact that $\mathcal{P}_r(e)$ and $\mathcal{P}_r(v)$ do not cross each other is easily checked using Fact 1. The second statement of the lemma follows from a few observations. Denote by v_1 the origin of e and by v_2 the end-vertex of e. If v is not on $\mathcal{P}_r(e)$ then clearly $\mathcal{P}_r(v)$ is not equal to $\mathcal{P}_r(e)$. If v is on $\mathcal{P}_r(e)$ between v_2 and N, then $\mathcal{P}_r(v) = \mathcal{P}_r(e)$ iff all edges of $\mathcal{P}_r(e)$ between v_1 and v are the rightmost ingoing red edge at their end-vertex. If v is on $\mathcal{P}_r(e)$ between S and v_1, then $\mathcal{P}_r(v) = \mathcal{P}_r(e)$ iff all edges of $\mathcal{P}_r(e)$ between v and v_2 are the leftmost outgoing red edge at their origin. It follows from these observations that $\mathcal{P}_r(v_2) = \mathcal{P}_r(e)$ if e is the rightmost ingoing red edge at its end-vertex, that $\mathcal{P}_r(v_1) = \mathcal{P}_r(e)$ if e is the leftmost outgoing red edge at its origin, and that no vertex v satisfies $\mathcal{P}_r(v) = \mathcal{P}_r(e)$ otherwise. $\square$

Definition. Given an irreducible triangulation T endowed with its minimal transversal edge-partition, a *ccw-internal* edge of T is an inner edge e of T such that the counterclockwise-consecutive edge at each extremity of e has the same color as e. Hence, on the minimal transversal pair of bipolar orientations, a ccw-internal red edge is an inner red edge of T that is neither the leftmost outgoing red edge at its origin nor the rightmost ingoing red edge at its end-vertex.

Lemma 13 *The number of abscissas not used by $\mathrm{DRAW}(T)$ is equal to the number of ccw-internal red edges of T. Similarly, the number of ordinates not used by $\mathrm{DRAW}(T)$ is equal to the number of ccw-internal blue edges of T.*

PROOF. Let $1 \leq i \leq f_r$ be an abscissa-candidate and let f be the unique inner face of T_r such that $\mathrm{Abs}(f) = i$. Recall that the separating red path $\mathcal{P}_r(f)$ is equal to the separating red path of the bottom-right edge e_f of f. Lemma 12 ensures that i is not the abscissa of any vertex of T iff e_f is a ccw-internal red edge of T. Hence, the number of abscissas not used by $\mathrm{DRAW}(T)$

is equal to the number of ccw-internal red edges of T that are the bottom-right edge of an inner face of T_r. This quantity is also the number of ccw-internal red edges of T. Indeed, a ccw-internal red edge e is the bottom-right edge of the inner face of T_r on its left, as e is not the leftmost outgoing red edge at its origin. □

6.2.3 Reduction to the analysis of a parameter on ternary trees

According to Lemma 13, the number of coordinate-lines that bear no vertex in $\mathrm{DRAW}(T)$ is equal to the number of ccw-internal edges of T. As we show next, the ccw-internal edges of T correspond to particular edges of the ternary tree A obtained by doing the opening of T. We define an *internal edge* of a ternary tree A as a closed edge e of A such that both edges following e in clockwise order around each extremity of e are closed edges.

Lemma 14 *Let T be an irreducible triangulation and let A be the ternary tree obtained by doing the opening of T, T being endowed with its minimal transversal edge-partition and A being endowed with the induced edge-bicoloration. Then each ccw-internal red (blue) edge of T corresponds, via the opening of T, to an internal red (blue, respectively) edge of A.*

PROOF. Let $e = (v, v')$ be a ccw-internal red edge of T. By definition, the ccw-consecutive edge of e at v' is red. Hence the cw-consecutive edge of e at v is blue (otherwise there would be a unicolored face in T, contradicting Lemma 1). Similarly, the cw-consecutive edge of e at v' is blue. Hence, by definition of the opening mapping, e is a closed edge of A. To prove that e is an internal edge of A, we have to prove that the cw-consecutive edge of A at each extremity of e is a closed edge of A. Let $e_1, \ldots, e_k$ be the clockwise interval of blue edges of T following e in clockwise order around v (the word interval refeering to the terminology of Condition C1); and let $v_1, \ldots, v_k$ be the corresponding sequence of neighbours of v. By definition of the opening mapping, the edge of A following e in clockwise order around v is e_k. To prove that e_k is a closed edge of A it remains to show that the edge $e' \in T$ following e_k in clockwise order around v_k is red. If $k = 1$, e' is also the edge following e in counterclockwise order around v'. Hence, the fact that e is internal red ensures that e' is red. If $k \geq 2$, then $e' = (v_k, v_{k-1})$. As (v, v_k) and (v, v_{k-1}) are blue, Lemma 1 ensures that (v_k, v_{k-1}) is red. Hence, e_k is a closed edge of A. Similarly, the edge of A following e in clockwise order around v' is a closed edge of A. Thus, e is an internal edge of A.

Conversely, let $e = (v, v')$ be an internal red edge of A. By definition of the opening mapping, the edge $e_v \in T$ following e in ccw order around v has the same color as e iff the half-edge of e_v incident to v has been created during
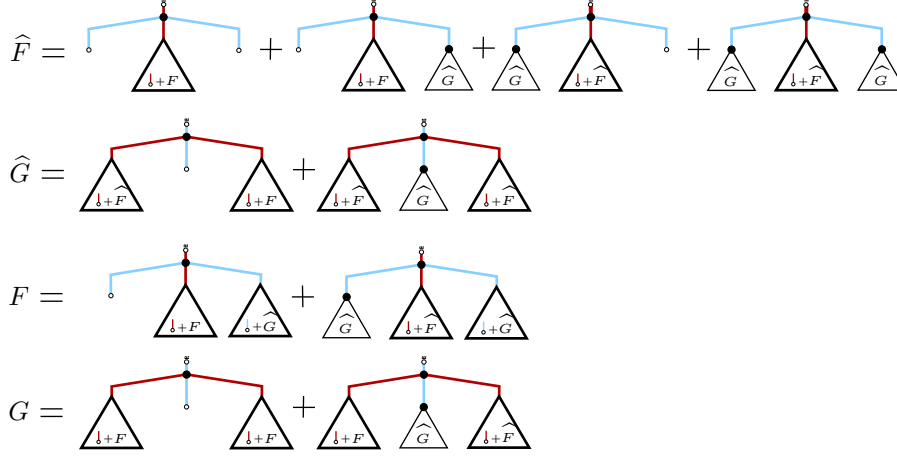
37

Fig. 13. Decomposition at the root node keeping track of the number of internal red edges.

a local closure. The same holds with the edge $e_{v'}$ following e in ccw order around v'. We claim that this condition is satisfied by e at v' (and similarly at v). Indeed, the sides of closed edges incident to the outer face of the figure F obtained as the partial closure of A are such that the cw-consecutive edge at their right extremity (looking toward the outer face) is a stem. Hence, e is not incident to the outer face of F, so that there exists a local closure whose effect is to close a triangular face incident to the right side of e (traversed from v to v'). The stem s involved in this local closure can not be incident to v, as the edge of A following e in cw order around v is not a stem. Hence $e = (v, v')$ is the second closed edge following the stem, so that the newly created half-edge (opposite to s) is incident to v'. This concludes the proof. □

Lemma 14 yields the following result that, as Fact 3, follows from an easy adaptation of the proof of Theorem 4.

Fact 4 *The distribution of the parameter $\Delta_{\mathrm{abs}}(T)$, for T taken uniformly at random in $\mathcal{T}'_n$, is equal to the distribution of the number of internal red edges in a rooted edge-bicolored ternary tree with n nodes taken uniformly at random.*

Let X_n be the random variable denoting the number of internal red edges of a rooted edge-bicolored ternary tree with n nodes taken uniformly at random. Fact 4 and the discussion in the overview ensure that the statement of Theorem 7 about W_c is proved by analyzing the distribution of X_n and showing a concentration around $5n/54$.

6.2.4 Analysis of the random variable X_n

We introduce the generating functions $F(z, u)$ and $G(z, u)$ counting respectively red-rooted ternary trees and blue-rooted ternary trees with respect to

38

the number of nodes and the number of internal red edges. As for the series $R(z, u)$ and $B(z, u)$ involved in the analysis of DRAW, an equation system linking the series $F(z, u)$ and $G(z, u)$ can be obtained by decomposing a rooted ternary tree at its root node into three subtrees.

To keep track of the parameter counting internal red edges, we introduce two auxiliary generating functions: $\widehat{F}(z, u)$ is the series counting red-rooted ternary trees with respect to the number of nodes and the number of internal red edges, with the difference that the root red stem is also considered as an internal red edge if its right child-edge is a closed edge; and $\widehat{G}(z, u)$ is the series counting blue-rooted ternary trees with respect to the number of nodes and the number of internal red edges, with the difference that the left child-edge e of the root node is also considered as an internal red edge if e is a closed edge and if the right child-edge of e is a closed edge. The decomposition at the root node keeping track of the number of internal red edges is shown in Figure 13. This directly translates to the following system:

$$
\begin{cases}
\widehat{F}(z, u) = z(1 + F) + zu(1 + F)\widehat{G} + z(1 + \widehat{F})\widehat{G} + zu\widehat{G}(1 + \widehat{F})\widehat{G} \\
\widehat{G}(z, u) = z(1 + \widehat{F})(1 + F) + z(1 + \widehat{F})\widehat{G}(1 + \widehat{F}) \\
F(z, u) = z(1 + F)(1 + \widehat{G}) + z\widehat{G}(1 + \widehat{F})(1 + \widehat{G}) \\
G(z, u) = z(1 + F)^2 + z(1 + F)\widehat{G}(1 + \widehat{F})
\end{cases}
\tag{3}
$$

Let $H(z, u)$ be the series counting rooted edge-bicolored ternary trees by the numbers of nodes and internal red edges; clearly $H(z, u) = F(z, u) + G(z, u)$. An algebraic equation $P(H(z, u), z, u) = 0$ is easily computed from (3). Then, the algebraic quasi-power theorem can be applied to the algebraic generating function $H(z, u)$. All conditions are easily checked and the algebraic series $\rho(u)$ which we obtain satisfies $\mu = -\rho'(1)/\rho(1) = 5/54$. This yields the statement of Theorem 7 on the distribution of W_c. The result is the same for H_c, because the distribution of H_c is equal to the distribution of W_c, by stability of $\mathcal{T}'_n$ and of COMPACTDRAW under the $\pi/2$ clockwise rotation. This concludes the proof of Theorem 7.

7   Conclusion

Several families of planar maps are characterised by the existence of a specific combinatorial structure for each map of the family (bipolar orientations for 2-connected maps, Schnyder woods for triangulations). This article provides a detailed study of so-called *transversal structures*, which are specific to the family of triangulations of the 4-gon with no filled 3-cycle —called irreducible

triangulations. The two main results we obtain related to transversal structures are a bijection between irreducible triangulations and ternary trees, as well as a new straight-line drawing algorithm for irreducible triangulations. The bijection allows us to count irreducible triangulations and the closely related family of 4-connected triangulations.

Surprisingly, the bijection —combined with some modern tools of analytic combinatorics— also makes it possible to analyse the grid size of the straight-line drawing; for a random irreducible triangulation with n vertices, the grid size is with high probability $11n/27 \times 11n/27$ up to an additive error $\mathcal{O}(\sqrt{n})$. In comparison, the best previously known algorithm [22] only guarantees a grid size $(\lceil n/2 \rceil - 1) \times \lfloor n/2 \rfloor$. The principle of our straight-line drawing algorithm has also been recently applied to the family of quadrangulations [15,16]; for a random quadrangulation with n vertices, the grid size is with high probability $13n/27 \times 13n/27$ up to an additive error $\mathcal{O}(\sqrt{n})$, which improves on the grid size $(\lceil n/2 \rceil - 1) \times \lfloor n/2 \rfloor$ guaranteed by [2].

References

[1] G. Di Battista, R. Tamassia, and I. G. Tollis. Area requirement and symmetry display of planar upward drawings. *Disc. Comput. Geometry*, 7:381–401, 1992.

[2] Therese Biedl and Franz J. Brandenburg. Drawing planar bipartite graphs with small area. In *Proceedings of CCCG, Windsor*, 2005.

[3] Therese C. Biedl, Goos Kant, and Michael Kaufmann. On triangulating planar graphs under the four-connectivity constraint. *Algorithmica*, 19(4):427–446, 1997.

[4] N. Bonichon, S. Felsner, and M. Mosbah. Convex drawings of 3-connected plane graphs. In *GD '04: Proceedings of the Symposium on Graph Drawing*, pages 287–299. Springer-Verlag, 2004.

[5] N. Bonichon, C. Gavoille, N. Hanusse, D. Poulalhon, and G. Schaeffer. Planar graphs, via well-orderly maps and trees. In 30^{th} *International Workshop, Graph - Theoretic Concepts in Computer Science (WG)*, volume 3353 of *Lecture Notes in Computer Science*, pages 270–284. Springer-Verlag, 2004.

[6] E. Brehm. 3-orientations and Schnyder 3-tree-decompositions. Master's thesis, FB Mathematik und Informatik, Freie Universität Berlin, 2000.

[7] W.G. Brown. Enumeration of triangulations of the disk. *Proceedings of the London Mathematical Society*, pages 746–768, 1964.

[8] M. de Berg, E. Mumford, and B. Speckmann. On rectilinear duals for vertex-weighted plane graphs. In *GD '05: Proceedings of the Symposium on Graph Drawing*. Springer-Verlag, 2006.

[9] H. de Fraysseix, P. Ossona de Mendez, and J. Pach. Representation of planar graphs by segments. *Intuitive Geometry*, 63:109–117, 1991.

[10] H. De Fraysseix, P. Ossona de Mendez, and P. Rosenstiehl. Bipolar orientations revisited. *Discrete Appl. Math.*, 56(2-3):157–179, 1995.

[11] Michael Drmota. Systems of functional equations. *Random Structures and Algorithms*, 10(1-2):103–124, 1997.

[12] S. Felsner. Convex drawings of planar graphs and the order dimension of 3-polytopes. *Order*, 18:19–37, 2001.

[13] S. Felsner. Lattice structures from planar graphs. *Electronic Journal of Combinatorics*, R15:24p., 2004.

[14] P. Flajolet and R. Sedgewick. Analytic combinatorics. Available at http://algo.inria.fr/flajolet/Publications/book051001.pdf. Preliminary version of the forthcoming book.

[15] E. Fusy. Straight-line drawing of quadrangulations. In *Proceedings of Graph Drawing'06, Karlsruhe*, pages 234–239. Springer, 2006.

[16] É. Fusy. *Combinatoire des cartes planaires et applications algorithmiques*. PhD thesis, École Polytechnique, 2007.

[17] É. Fusy, D. Poulalhon, and G. Schaeffer. Dissections and trees, with applications to optimal mesh encoding and to random sampling. In *16th Annual ACM-SIAM Symposium on Discrete Algorithms*, January 2005.

[18] X. He. On finding the rectangular duals of planar triangulated graphs. *SIAM J. Comput.*, 22:1218–1226, 1993.

[19] X. He. Grid embedding of 4-connected plane graphs. *Discrete Computational Geometry*, pages 339–358, 1997.

[20] G. Kant and Xin He. Regular edge labeling of 4-connected plane graphs and its applications in graph drawing problems. *Theoretical Computer Science*, 172(1-2):175–193, 1997.

[21] S. Khuller, J. Naor, and P. N. Klein. The lattice structure of flow in planar graphs. *SIAM J. Discrete Math.*, 6(3):477–490, 1993.

[22] K. Miura, S. Nakano, and T. Nishizeki. Grid drawings of four-connected plane graphs. *Disc. Comput. Geometry*, 26(2):73–87, 2001.

[23] P. Ossona de Mendez. *Orientations bipolaires*. PhD thesis, Paris, 1994.

[24] D. Poulalhon and G. Schaeffer. Optimal coding and sampling of triangulations. In *Automata, Languages and Programming. 30th International Colloquium, ICALP 2003, Eindhoven*, volume 2719, pages 1080–1094, 2003.

[25] G. Schaeffer. *Conjugaison d'arbres et cartes combinatoires aléatoires*. PhD thesis, Université Bordeaux I, 1998.

[26] Gilles Schaeffer. Random sampling of large planar maps and convex polyhedra. In *Annual ACM Symposium on Theory of Computing (Atlanta, GA, 1999)*, pages 760–769 (electronic). ACM, New York, 1999.

[27] W. Schnyder. Planar graphs and poset dimension. *Order*, 5:323–343, 1989.

[28] W. Schnyder. Embedding planar graphs on the grid. In *SODA '90: Proceedings of the first annual ACM-SIAM symposium on Discrete algorithms*, pages 138–148, 1990.

[29] B. Speckmann and M. Van Kreveld. On rectangular cartograms, 2004. Technical report.

[30] W. T. Tutte. A census of planar triangulations. *Canad. J. Math.*, 14:21–38, 1962.